

Low Power Dynamic Scheduling for Computing Systems

Michael J. Neely

Abstract

This paper considers energy-aware control for a computing system with two states: *active* and *idle*. In the active state, the controller chooses to perform a single task using one of multiple task processing modes. The controller then saves energy by choosing an amount of time for the system to be idle. These decisions affect processing time, energy expenditure, and an abstract *attribute vector* that can be used to model other criteria of interest (such as processing quality or distortion). The goal is to optimize time average system performance. Applications of this model include a smart phone that makes energy-efficient computation and transmission decisions, a computer that processes tasks subject to rate, quality, and power constraints, and a smart grid energy manager that allocates resources in reaction to a time varying energy price. The solution methodology of this paper uses the theory of *optimization for renewal systems* developed in our previous work. This paper is written in tutorial form and develops the main concepts of the theory using several detailed examples. It also highlights the relationship between online dynamic optimization and linear fractional programming. Finally, it provides exercises to help the reader learn the main concepts and apply them to their own optimizations. This paper is an arxiv technical report, and is a preliminary version of material that will appear as a book chapter in an upcoming book on green communications and networking.

Index Terms

Queueing analysis, optimization, stochastic control, renewal theory

I. INTRODUCTION

This paper considers energy-aware control for a computing system with two states: *active* and *idle*. In the active state, the controller chooses to perform a single task using one of multiple task processing modes. The controller then saves energy by choosing an amount of time for the system to be idle. These decisions affect processing time, energy expenditure, and an abstract *attribute vector* that can be used to model other criteria of interest (such as processing quality or distortion). The goal is to optimize time average system performance. Applications of this model include a smart phone that makes energy-efficient computation and transmission decisions, a computer that processes tasks subject to rate, quality, and power constraints, and a smart grid energy manager that allocates resources in reaction to a time varying energy price.

The solution methodology of this paper uses the theory of *optimization for renewal systems* developed in [1][2]. Section II focuses on a computer system that seeks to minimize average power subject to processing rate constraints for different classes of tasks. Section III generalizes to treat optimization for a larger class of systems. Section IV extends the model to allow control actions to react to a random event observed at the beginning of each active period, such as a vector of current channel conditions or energy prices.

II. TASK SCHEDULING WITH PROCESSING RATE CONSTRAINTS

To illustrate the method, this section considers a particular system. Consider a computer system that repeatedly processes tasks. There are N classes of tasks, where N is a positive integer. For simplicity, assume each class always has a new task ready to be performed (this is extended to randomly arriving tasks in Section III-C). The system operates over time intervals called *frames*. Each frame $k \in \{0, 1, 2, \dots\}$ begins with an *active period* of size $D[k]$ and ends with an *idle period* of size $I[k]$ (see Fig. 1). At the beginning of each active period k , the controller selects a new task of class $c[k] \in \{1, \dots, N\}$. It also chooses a *processing mode* $m[k]$ from a finite set

The author is with the Electrical Engineering department at the University of Southern California, Los Angeles, CA.

This material is supported in part by one or more of the following: the NSF Career grant CCF-0747525, the Network Science Collaborative Technology Alliance sponsored by the U.S. Army Research Laboratory W911NF-09-2-0053.

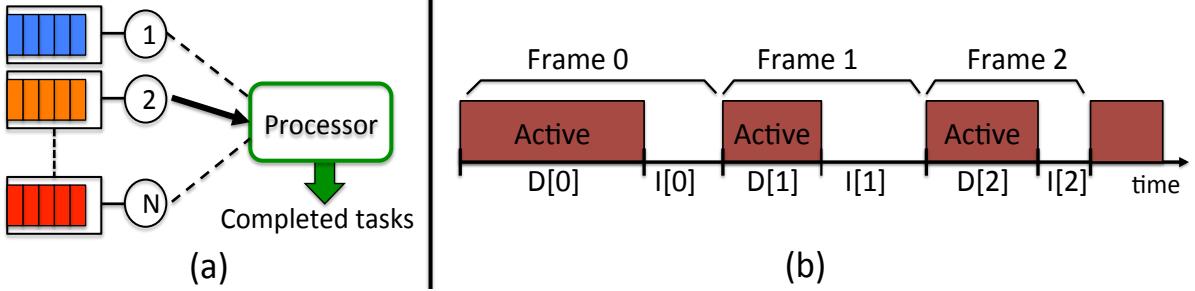


Fig. 1. (a) A processor that chooses from one of N classes on each frame k . (b) A timeline illustrating the active and idle periods for each frame.

\mathcal{M} of possible processing options. These control decisions affect the duration $D[k]$ of the active period and the energy $e[k]$ that is incurred. The controller then selects an amount of time $I[k]$ to be idle, where $I[k]$ is chosen such that $0 \leq I[k] \leq I_{max}$ for some positive number I_{max} . Choosing $I[k] = 0$ effectively skips the idle period, so there can be back-to-back active periods. For simplicity, this section assumes that no energy is expended in the idle state.

Assume that $D[k]$ and $e[k]$ are *random functions* of the class and mode decisions for frame k . Specifically, assume $D[k]$ and $e[k]$ are conditionally independent of the past, given the current $(c[k], m[k])$ that is used, with mean values given by functions $\hat{D}(c, m)$ and $\hat{e}(c, m)$ defined over $(c, m) \in \{1, \dots, N\} \times \mathcal{M}$:

$$\hat{D}(c, m) \triangleq \mathbb{E}[D[k]|(c[k], m[k]) = (c, m)] , \quad \hat{e}(c, m) \triangleq \mathbb{E}[e[k]|(c[k], m[k]) = (c, m)]$$

where the notation “ $a \triangleq b$ ” represents “ a is defined to be equal to b .” This paper uses $\hat{D}(c[k], m[k])$ and $\hat{e}(c[k], m[k])$ to denote expectations given a particular decision $(c[k], m[k])$ for frame k :

$$\hat{D}(c[k], m[k]) = \mathbb{E}[D[k]|c[k], m[k]] , \quad \hat{e}(c[k], m[k]) = \mathbb{E}[e[k]|c[k], m[k]]$$

It is assumed there is a positive value D_{min} such that $D[k] \geq D_{min}$ for all frames k , regardless of the $(c[k], m[k])$ decisions. Thus, all frame sizes are at least D_{min} units of time. Further, for technical reasons, it is assumed that second moments of $D[k]$ and $e[k]$ are bounded by a finite constant σ^2 , so that:

$$\mathbb{E}[D[k]^2] \leq \sigma^2 , \quad \mathbb{E}[e[k]^2] \leq \sigma^2 \quad (1)$$

where (1) holds regardless of the policy for selecting $(c[k], m[k])$. The conditional joint distribution of $(D[k], e[k])$, given $(c[k], m[k])$, is otherwise arbitrary, and only the mean values $\hat{D}(c, m)$ and $\hat{e}(c, m)$ are known for each $c \in \{1, \dots, N\}$ and $m \in \mathcal{M}$. In the special case of a deterministic system, the functions $\hat{D}(c, m)$ and $\hat{e}(c, m)$ can be viewed as deterministic mappings from a given control action $(c[k], m[k]) = (c, m)$ to the actual delay $D[k] = \hat{D}(c, m)$ and energy $e[k] = \hat{e}(c, m)$ experienced on frame k , rather than as expectations of these values.

A. Examples of Energy-Aware Processing

Consider an example where a computer system performs computation for different tasks. Suppose the system uses a chip multi-processor that can select between one of multiple processing modes for each task. For example, this might be done using voltage/frequency scaling, or by using a choice of different processing cores [3][4]. Let \mathcal{M} represent the set of processing modes. For each mode $m \in \mathcal{M}$, define:

$$\begin{aligned} T_{\text{setup}}(m) &\triangleq \text{Setup time for mode } m. \\ e_{\text{setup}}(m) &\triangleq \text{Setup energy for mode } m. \\ \text{DPI}(m) &\triangleq \text{Average delay-per-instruction for mode } m. \\ \text{EPI}(m) &\triangleq \text{Average energy-per-instruction for mode } m. \end{aligned}$$

Further suppose that tasks of class $c \in \{1, \dots, N\}$ have an average number of instructions equal to \bar{S}_c . For simplicity, suppose the number of instructions in a task is independent of the energy and delay of each individual

instruction. Then the average energy and delay functions $\hat{e}(c, m)$ and $\hat{D}(c, m)$ are:

$$\begin{aligned}\hat{e}(c, m) &= e_{\text{setup}}(m) + \bar{S}_c EPI(m) \\ \hat{D}(c, m) &= T_{\text{setup}}(m) + \bar{S}_c DPI(m)\end{aligned}$$

In cases when the system cannot be modeled using $DPI(m)$ and $EPI(m)$ functions, the expectations $\hat{e}(c, m)$ and $\hat{D}(c, m)$ can be estimated as empirical averages of energy and delay observed when processing type c tasks with mode m . While this example assumes all classes have the same set of processing mode options \mathcal{M} , this can easily be extended to restrict each class c to its own subset of options \mathcal{M}_c .

As another example, consider the problem of *wireless data transmission*. Here, each task represents a packet of data that must be transmitted. Let \mathcal{M} represent the set of wireless transmission options (such as modulation and coding strategies). For each $m \in \mathcal{M}$, define $\mu(m)$ as the transmission rate (in bits per unit time) under option m , and let $power(m)$ be the power used. For simplicity, assume there are no transmission errors. Let \bar{B}_c represent the average packet size for class $c \in \{1, \dots, N\}$, in units of bits. Thus:

$$\begin{aligned}\hat{e}(c, m) &= power(m)\bar{B}_c/\mu(m) \\ \hat{D}(c, m) &= \bar{B}_c/\mu(m)\end{aligned}$$

In the case when each transmission mode $m \in \mathcal{M}$ has a known error probability, the functions $\hat{e}(c, m)$ and $\hat{D}(c, m)$ can be redefined to account for retransmissions. An important alternative scenario is when channel states are time-varying but can be measured at the beginning of each frame. This can be treated using the extended theory in Section IV.

B. Time Averages as Ratios of Frame Averages

The goal is to design a control policy that makes decisions over frames to minimize time average power subject to processing each class $n \in \{1, \dots, N\}$ with rate at least λ_n , for some desired processing rates $(\lambda_1, \dots, \lambda_N)$ that are given. Before formalizing this as a mathematical optimization, this subsection shows how to write time averages in terms of frame averages. Suppose there is a control policy that yields a sequence of energies $\{e[0], e[1], e[2], \dots\}$ and corresponding frame sizes $\{D[0] + I[0], D[1] + I[1], D[2] + I[2], \dots\}$ for each frame $k \in \{0, 1, 2, \dots\}$. The *frame averages* \bar{e} , \bar{D} , \bar{I} are defined:

$$\bar{e} \triangleq \lim_{K \rightarrow \infty} \frac{1}{K} \sum_{k=0}^{K-1} e[k] , \quad \bar{D} \triangleq \lim_{K \rightarrow \infty} \frac{1}{K} \sum_{k=0}^{K-1} D[k] , \quad \bar{I} \triangleq \lim_{K \rightarrow \infty} \frac{1}{K} \sum_{k=0}^{K-1} I[k] \quad (2)$$

where, for simplicity, it is assumed the limits converge to constants with probability 1. Note that \bar{e} does *not* represent the time average power used by the system, because it does not consider the amount of time spent in each frame. The time average power considers the accumulated energy used divided by the total time, and is written as follows:

$$\lim_{K \rightarrow \infty} \frac{\sum_{k=0}^{K-1} e[k]}{\sum_{k=0}^{K-1} (D[k] + I[k])} = \lim_{K \rightarrow \infty} \frac{\frac{1}{K} \sum_{k=0}^{K-1} e[k]}{\frac{1}{K} \sum_{k=0}^{K-1} (D[k] + I[k])} = \frac{\bar{e}}{\bar{D} + \bar{I}}$$

Therefore, the time average power is equal to the average energy per frame divided by the average frame size. This simple observation is often used in *renewal-reward theory* [5][6].

For each class $n \in \{1, \dots, N\}$ and each frame k , define an *indicator variable* $1_n[k]$ that is 1 if the controller chooses to process a class n task on frame k , and 0 else:

$$1_n[k] \triangleq \begin{cases} 1 & \text{if } c[k] = n \\ 0 & \text{if } c[k] \neq n \end{cases}$$

Then $\bar{1}_n \triangleq \lim_{K \rightarrow \infty} \frac{1}{K} \sum_{k=0}^{K-1} 1_n[k]$ is the fraction of frames that choose class n , and the ratio $\bar{1}_n / (\bar{D} + \bar{I})$ is the time average rate of processing class n tasks, in tasks per unit time.

The problem of minimizing time average power subject to processing each class n at a rate of at least λ_n tasks per unit time is then mathematically written as follows:

$$\text{Minimize: } \frac{\bar{e}}{\bar{D} + \bar{I}} \quad (3)$$

$$\text{Subject to: } \frac{\bar{I}_n}{\bar{D} + \bar{I}} \geq \lambda_n \quad \forall n \in \{1, \dots, N\} \quad (4)$$

$$(c[k], m[k]) \in \{1, \dots, N\} \times \mathcal{M} \quad \forall k \in \{0, 1, 2, \dots\} \quad (5)$$

$$0 \leq I[k] \leq I_{max} \quad \forall k \in \{0, 1, 2, \dots\} \quad (6)$$

where the objective (3) is average power, the constraint (4) ensures the processing rate of each class n is at least λ_n , and constraints (5)-(6) ensure that $c[k] \in \{1, \dots, N\}$, $m[k] \in \mathcal{M}$, and $0 \leq I[k] \leq I_{max}$ for each frame k .

C. Relation to Frame Average Expectations

The problem (3)-(6) is defined by frame averages. This subsection shows that frame averages are related to frame average *expectations*, and hence can be related to the expectation functions $\hat{D}(c, m)$, $\hat{e}(c, m)$. Consider any (possibly randomized) control algorithm for selecting $(c[k], m[k])$ over frames, and assume this gives rise to well defined expectations $\mathbb{E}[e[k]]$ for each frame k . By the law of iterated expectations, it follows that for any given frame $k \in \{0, 1, 2, \dots\}$:

$$\mathbb{E}[e[k]] = \mathbb{E}[\mathbb{E}[e[k]|c[k], m[k]]] = \mathbb{E}[\hat{e}(c[k], m[k])] \quad (7)$$

Furthermore, because second moments are bounded by a constant σ^2 on each frame k , the *bounded moment convergence theorem* (given in Appendix A) ensures that if the frame average energy converges to a constant \bar{e} with probability 1, as defined in (2), then \bar{e} is the same as the frame average expectation:

$$\lim_{K \rightarrow \infty} \frac{1}{K} \sum_{k=0}^{K-1} e[k] = \bar{e} = \lim_{K \rightarrow \infty} \frac{1}{K} \sum_{k=0}^{K-1} \mathbb{E}[\hat{e}(c[k], m[k])]$$

The same goes for the quantities \bar{D} , \bar{I} , \bar{I}_n . Hence, one can interpret the problem (3)-(6) using frame average expectations, rather than pure frame averages.

D. An Example with One Task Class

Consider a simple example with only one type of task. The system processes a new task of this type at the beginning of every busy period. The energy and delay functions can then be written purely in terms of the processing mode $m \in \mathcal{M}$, so that we have $\hat{e}(m)$ and $\hat{D}(m)$. Suppose there are only two processing mode options, so that $\mathcal{M} = \{1, 2\}$, and that each option leads to a deterministic energy and delay, as given below:

$$m[k] = 1 \implies (e[k], D[k]) = (\hat{e}(1), \hat{D}(1)) = (1, 7) \quad (8)$$

$$m[k] = 2 \implies (e[k], D[k]) = (\hat{e}(2), \hat{D}(2)) = (3, 4) \quad (9)$$

Option $m[k] = 1$ requires 1 unit of energy but 7 units of processing time. Option $m[k] = 2$ is more energy-expensive (requiring 3 units of energy) but is faster (taking only 4 time units). The idle time $I[k]$ is chosen every frame in the interval $[0, 10]$, so that $I_{max} = 10$.

1) *No constraints*: For this system, suppose we seek to minimize average power $\bar{e}/(\bar{D} + \bar{I})$, with no processing rate constraint. Consider three possible algorithms:

- 1) Always use $m[k] = 1$ and $I[k] = I$ for all frames k , for some constant $I \in [0, 10]$.
- 2) Always use $m[k] = 2$ and $I[k] = I$ for all frames k , for some constant $I \in [0, 10]$.
- 3) Always use $I[k] = I$ for all frames k , for some constant $I \in [0, 10]$. However, each frame k , independently choose $m[k] = 1$ with probability p , and $m[k] = 2$ with probability $1 - p$ (for some p that satisfies $0 \leq p \leq 1$).

Clearly the third algorithm contains the first two for $p = 1$ and $p = 0$, respectively. The time average power under each algorithm is:

$$\begin{aligned} m[k] = 1 \text{ always} &\implies \frac{\bar{e}}{\bar{D} + \bar{I}} = \frac{\hat{e}(1)}{\hat{D}(1) + I} = \frac{1}{7 + I} \\ m[k] = 2 \text{ always} &\implies \frac{\bar{e}}{\bar{D} + \bar{I}} = \frac{\hat{e}(2)}{\hat{D}(2) + I} = \frac{3}{4 + I} \\ \text{probabilistic rule} &\implies \frac{\bar{e}}{\bar{D} + \bar{I}} = \frac{p\hat{e}(1) + (1-p)\hat{e}(2)}{p\hat{D}(1) + (1-p)\hat{D}(2) + I} = \frac{1p + 3(1-p)}{7p + 4(1-p) + I} \end{aligned}$$

It is clear that in all three cases, we should choose $I = 10$ to minimize average power. Further, it is clear that always choosing $m[k] = 1$ is better than always choosing $m[k] = 2$. However, it is not immediately obvious if a randomized mode selection rule can do even better. The answer is no: In this case, power is minimized by choosing $m[k] = 1$ and $I[k] = 10$ for all frames k , yielding average power $1/17$.

This fact holds more generally: Let $a(\alpha)$ and $b(\alpha)$ be any deterministic real valued functions defined over general *actions* α that are chosen in an abstract *action space* \mathcal{A} . Assume the functions are bounded, and that there is a value $b_{min} > 0$ such that $b(\alpha) \geq b_{min}$ for all $\alpha \in \mathcal{A}$. Consider designing a randomized choice of α that minimizes $\mathbb{E}[a(\alpha)] / \mathbb{E}[b(\alpha)]$, where the expectations are with respect to the randomness in the α selection. The next lemma shows this is done by *deterministically* choosing an action $\alpha \in \mathcal{A}$ that minimizes $a(\alpha)/b(\alpha)$.

Lemma 1: Under the assumptions of the preceding paragraph, for any randomized choice of $\alpha \in \mathcal{A}$ we have:

$$\frac{\mathbb{E}[a(\alpha)]}{\mathbb{E}[b(\alpha)]} \geq \inf_{\alpha \in \mathcal{A}} \left[\frac{a(\alpha)}{b(\alpha)} \right]$$

and so the infimum of the ratio over the class of deterministic decisions yields a value that is less than or equal to that of any randomized selection.

Proof: Consider any randomized policy that yields expectations $\mathbb{E}[a(\alpha)]$ and $\mathbb{E}[b(\alpha)]$. Without loss of generality, assume these expectations are achieved by a policy that randomizes over a finite set of M actions $\alpha_1, \dots, \alpha_M$ in \mathcal{A} with some probabilities p_1, \dots, p_M :¹

$$\mathbb{E}[a(\alpha)] = \sum_{m=1}^M p_m a(\alpha_m), \quad \mathbb{E}[b(\alpha)] = \sum_{m=1}^M p_m b(\alpha_m)$$

Then, because $b(\alpha_m) > 0$ for all α_m , we have:

$$\begin{aligned} \frac{\mathbb{E}[a(\alpha)]}{\mathbb{E}[b(\alpha)]} &= \frac{\sum_{m=1}^M p_m a(\alpha_m)}{\sum_{m=1}^M p_m b(\alpha_m)} \\ &= \frac{\sum_{m=1}^M p_m b(\alpha_m) [a(\alpha_m)/b(\alpha_m)]}{\sum_{m=1}^M p_m b(\alpha_m)} \\ &\geq \frac{\sum_{m=1}^M p_m b(\alpha_m) \inf_{\alpha \in \mathcal{A}} [a(\alpha)/b(\alpha)]}{\sum_{m=1}^M p_m b(\alpha_m)} \\ &= \inf_{\alpha \in \mathcal{A}} [a(\alpha)/b(\alpha)] \end{aligned}$$

□

2) *One Constraint:* The preceding subsection shows that unconstrained problems can be solved by deterministic actions. This is not true for constrained problems. This subsection shows that adding just a single constraint often *necessitates* the use of randomized actions. Consider the same problem as before, with two choices for $m[k]$, and with the same $\hat{e}(m)$ and $\hat{D}(m)$ values given in (8)-(9). We want to minimize $\bar{e}/(\bar{D} + \bar{I})$ subject to $1/(\bar{D} + \bar{I}) \geq 1/5$, where $1/(\bar{D} + \bar{I})$ is the rate of processing jobs. The constraint is equivalent to $\bar{D} + \bar{I} \leq 5$.

¹Indeed, because the set $\mathcal{S} = \{(a(\alpha), b(\alpha)) \text{ such that } \alpha \in \mathcal{A}\}$ is bounded, the expectation $(\mathbb{E}[a(\alpha)], \mathbb{E}[b(\alpha)])$ is finite and is contained in the convex hull of \mathcal{S} . Thus, $(\mathbb{E}[a(\alpha)], \mathbb{E}[b(\alpha)])$ is a convex combination of a finite number of points in \mathcal{S} .

Assume the algorithm chooses $I[k]$ over frames to yield an average \bar{I} that is somewhere in the interval $0 \leq \bar{I} \leq 10$. Now consider different algorithms for selecting $m[k]$. If we choose $m[k] = 1$ always (so that $(\hat{e}(1), \hat{D}(1)) = (1, 7)$), then:

$$m[k] = 1 \text{ always} \implies \frac{\bar{e}}{\bar{D} + \bar{I}} = \frac{1}{7 + \bar{I}} , \quad \bar{D} + \bar{I} = 7 + \bar{I}$$

and so it is impossible to meet the constraint $\bar{D} + \bar{I} \leq 5$, because $7 + \bar{I} > 5$.

If we choose $m[k] = 2$ always (so that $(\hat{e}(2), \hat{D}(2)) = (3, 4)$), then:

$$m[k] = 2 \text{ always} \implies \frac{\bar{e}}{\bar{D} + \bar{I}} = \frac{3}{4 + \bar{I}} , \quad \bar{D} + \bar{I} = 4 + \bar{I}$$

It is clear that we can meet the constraint by choosing \bar{I} so that $0 \leq \bar{I} \leq 1$, and power is minimized in this setting by using $\bar{I} = 1$. This can be achieved, for example, by using $I[k] = 1$ for all frames k . This meets the processing rate constraint with equality: $\bar{D} + \bar{I} = 4 + 1 = 5$. Further, it yields average power $\bar{e}/(\bar{D} + \bar{I}) = 3/5 = 0.6$.

However, it is possible to reduce average power while also meeting the constraint with equality by using the following randomized policy (which can be shown to be optimal): Choose $I[k] = 0$ for all frames k , so that $\bar{I} = 0$. Then every frame k , independently choose $m[k] = 1$ with probability $1/3$, and $m[k] = 2$ with probability $2/3$. We then have:

$$\bar{D} + \bar{I} = (1/3)7 + (2/3)4 + 0 = 5$$

and so the processing rate constraint is met with equality. However, average power is:

$$\frac{\bar{e}}{\bar{D} + \bar{I}} = \frac{(1/3)1 + (2/3)3}{(1/3)7 + (2/3)4 + 0} = 7/15 \approx 0.466667$$

This is a significant savings over the average power of 0.6 from the deterministic policy.

E. The Linear Fractional Program for Task Scheduling

Now consider the general problem (3)-(6) for minimizing average power subject to average processing rate constraints for each of the L classes. Assume the problem is *feasible*, so that it is possible to choose actions $(c[k], m[k], I[k])$ over frames to meet the desired constraints (4)-(6). It can be shown that an optimal solution can be achieved over the class of *stationary and randomized policies* with the following structure: Every frame, independently choose vector $(c[k], m[k])$ with some probabilities $p(c, m) = \Pr[(c[k], m[k]) = (c, m)]$. Further, use a constant idle time $I[k] = I$ for all frames k , for some constant I that satisfies $0 \leq I \leq I_{max}$. Thus, the problem can be written as the following *linear fractional program* with unknowns $p(c, m)$ and I and known constants $\hat{e}(c, m)$, $\hat{D}(c, m)$, λ_n , and I_{max} :

$$\text{Minimize: } \frac{\sum_{c=1}^N \sum_{m \in \mathcal{M}} p(c, m) \hat{e}(c, m)}{I + \sum_{c=1}^N \sum_{m \in \mathcal{M}} p(c, m) \hat{D}(c, m)} \quad (10)$$

$$\text{Subject to: } \frac{\sum_{m \in \mathcal{M}} p(n, m)}{I + \sum_{c=1}^N \sum_{m \in \mathcal{M}} p(c, m) \hat{D}(c, m)} \geq \lambda_n \quad \forall n \in \{1, \dots, N\} \quad (11)$$

$$0 \leq I \leq I_{max} \quad (12)$$

$$p(c, m) \geq 0 \quad \forall c \in \{1, \dots, N\}, \quad \forall m \in \mathcal{M} \quad (13)$$

$$\sum_{c=1}^N \sum_{m \in \mathcal{M}} p(c, m) = 1 \quad (14)$$

where the numerator and denominator in (10) are equal to \bar{e} and $\bar{D} + \bar{I}$, respectively, under this randomized algorithm, the numerator in the left-hand-side of (11) is equal to \bar{I}_n , and the constraints (13)-(14) specify that $p(c, m)$ must be a valid probability mass function.

Linear fractional programs can be solved in several ways. One method uses a nonlinear change of variables to map the problem to a convex program [7]. However, this method does not admit an *online* implementation, because time averages are not preserved through the nonlinear change of variables. Below, an online algorithm is presented that makes decisions every frame k . The algorithm is *not* a stationary and randomized algorithm as described above.

However, it yields time averages that satisfy the desired constraints of the problem (3)-(6), with a time average power expenditure that can be pushed arbitrarily close to the optimal value. A significant advantage of this approach is that it extends to treat cases with random task arrivals, without requiring knowledge of the $(\lambda_1, \dots, \lambda_N)$ arrival rates, and to treat other problems with observed random events, without requiring knowledge of the probability distribution for these events. These extensions are shown in later sections.

For later analysis, it is useful to write (10)-(14) in a simpler form. Let $power^{opt}$ be the optimal time average power for the above linear fractional program, achieved by some probability distribution $p^*(c, m)$ and idle time I^* that satisfies $0 \leq I^* \leq I_{max}$. Let $(c^*[k], m^*[k], I^*[k])$ represent the frame k decisions under this stationary and randomized policy. Thus:

$$\frac{\mathbb{E}[\hat{e}(c^*[k], m^*[k])]}{\mathbb{E}[I^*[k] + \hat{D}(c^*[k], m^*[k])]} = power^{opt} \quad (15)$$

$$\frac{\mathbb{E}[1_n^*[k]]}{\mathbb{E}[I^*[k] + \hat{D}(c^*[k], m^*[k])]} \geq \lambda_n \quad \forall n \in \{1, \dots, N\} \quad (16)$$

where $1_n^*[k]$ is an indicator function that is 1 if $c^*[k] = n$, and 0 else. The numerator and denominator of (15) correspond to those of (10). Likewise, the constraint (16) corresponds to (11).

F. Virtual Queues

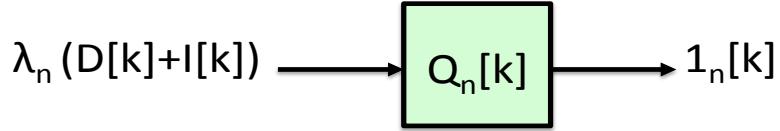


Fig. 2. An illustration of the virtual queue $Q_n[k]$ from equation (18).

To solve the problem (3)-(6), we first consider the constraints (4), which are equivalent to the constraints:

$$\lambda_n(\bar{D} + \bar{I}) \leq \bar{1}_n \quad \forall n \in \{1, \dots, N\} \quad (17)$$

For each constraint $n \in \{1, \dots, N\}$, define a *virtual queue* $Q_n[k]$ that is updated on frames $k \in \{0, 1, 2, \dots\}$ by:

$$Q_n[k+1] = \max[Q_n[k] + \lambda_n(D[k] + I[k]) - 1_n[k], 0] \quad (18)$$

The initial condition $Q_n[0]$ can be any non-negative value. For simplicity, it is assumed throughout that $Q_n[0] = 0$ for all $n \in \{1, \dots, N\}$. The update (18) can be viewed as a discrete time queueing equation, where $Q_n[k]$ is the backlog on frame k , $\lambda_n(D[k] + I[k])$ is an effective amount of “new arrivals,” and $1_n[k]$ is the amount of “offered service” (see Fig. 2). The intuition is that if all virtual queues $Q_n[k]$ are *stable*, then the average “arrival rate” $\lambda_n(\bar{D} + \bar{I})$ must be less than or equal to the average “service rate” $\bar{1}_n$, which ensures the desired constraint (17). This is made precise in the following lemma.

Lemma 2: (Virtual Queues) Suppose $Q_n[k]$ has update equation given by (18), with any non-negative initial condition.

a) For all $K \in \{1, 2, 3, \dots\}$ we have:

$$\frac{1}{K} \sum_{k=0}^{K-1} [\lambda_n(D[k] + I[k]) - 1_n[k]] \leq \frac{Q_n[K] - Q_n[0]}{K} \quad (19)$$

b) If $\lim_{K \rightarrow \infty} Q_n[K]/K = 0$ with probability 1, then:

$$\limsup_{K \rightarrow \infty} \frac{1}{K} \sum_{k=0}^{K-1} [\lambda_n(D[k] + I[k]) - 1_n[k]] \leq 0 \quad \text{with probability 1} \quad (20)$$

c) If $\lim_{K \rightarrow \infty} \mathbb{E}[Q_n[K]]/K = 0$, then:

$$\limsup_{K \rightarrow \infty} [\lambda_n(\bar{D}[K] + \bar{I}[K]) - \bar{1}_n[K]] \leq 0 \quad (21)$$

where $\bar{D}[K]$, $\bar{I}[K]$, $\bar{1}_n[K]$ are defined:

$$\bar{D}[K] \triangleq \frac{1}{K} \sum_{k=0}^{K-1} \mathbb{E}[D[k]] \quad , \quad \bar{I}[K] \triangleq \frac{1}{K} \sum_{k=0}^{K-1} \mathbb{E}[I[k]] \quad , \quad \bar{1}_n[K] \triangleq \frac{1}{K} \sum_{k=0}^{K-1} \mathbb{E}[1_n[k]]$$

Proof: From (18) we have for all $k \in \{0, 1, 2, \dots\}$:

$$Q_n[k+1] \geq Q_n[k] + \lambda_n(D[k] + I[k]) - 1_n[k]$$

Fixing a positive integer K and summing the above over $k \in \{0, \dots, K-1\}$ yields:

$$Q_n[K] - Q_n[0] \geq \sum_{k=0}^{K-1} [\lambda_n(D[k] + I[k]) - 1_n[k]]$$

Dividing the above by K proves part (a). Part (b) follows from (19) by taking a \limsup . Part (c) follows by first taking expectations of (19) and then taking a \limsup . \square

Inequality (19) shows that the value $Q_n[K]/K$ bounds the amount by which the desired constraint for class n is violated by the time averages achieved over the first K frames. Suppose that $D[k]$, $I[k]$, and $1_n[k]$ have frame averages that converge to constants \bar{D} , \bar{I} , $\bar{1}_n$ with probability 1. Part (c) of Lemma 2 indicates that if $\lim_{K \rightarrow \infty} \mathbb{E}[Q_n[K]]/K = 0$ for all $n \in \{1, \dots, N\}$, then $\lambda_n(\bar{D} + \bar{I}) \leq \bar{1}_n$ for all $n \in \{1, \dots, N\}$.

In the language of queueing theory, a discrete time queue $Q[k]$ is said to be *rate stable* if $\lim_{k \rightarrow \infty} Q[k]/k = 0$ with probability 1, and is *mean rate stable* if $\lim_{k \rightarrow \infty} \mathbb{E}[Q[k]]/k = 0$ [1]. With this terminology, the above lemma shows that if $Q_n[k]$ is rate stable then (20) holds, and if $Q_n[k]$ is mean rate stable then (21) holds.

G. The Drift-Plus-Penalty Ratio

To stabilize the queues while minimizing time average power, we use *Lyapunov optimization theory*, which gives rise to the *drift-plus-penalty ratio algorithm* [1]. First define $L[k]$ as the sum of the squares of all queues on frame k (divided by 2 for convenience later):

$$L[k] \triangleq \frac{1}{2} \sum_{n=1}^N Q_n[k]^2$$

$L[k]$ is often called a *Lyapunov function*, and acts as a scalar measure of the size of the queues. Intuitively, keeping $L[k]$ small leads to stable queues, and we should take actions that tend to shrink $L[k]$ from one frame to the next. Define $\Delta[k]$ as the *Lyapunov drift*, being the difference in the Lyapunov function from one frame to the next:

$$\Delta[k] \triangleq L[k+1] - L[k]$$

Taking actions to minimize $\Delta[k]$ every frame can be shown to ensure the desired constraints are satisfied whenever it is possible to satisfy them, but does not incorporate power minimization. To incorporate this, every frame k we observe the current queue vector $\mathbf{Q}[k] = (Q_1[k], \dots, Q_N[k])$ and choose control actions $(c[k], m[k], I[k])$ to minimize a bound on the following *drift-plus-penalty ratio*:

$$\frac{\mathbb{E}[\Delta[k] + V e[k] | \mathbf{Q}[k]]}{\mathbb{E}[D[k] + I[k] | \mathbf{Q}[k]]}$$

where V is a non-negative parameter that weights the extent to which power minimization is emphasized. The intuition is that the numerator incorporates both drift and energy. The denominator “normalizes” this by the expected frame size, with the understanding that average power must include both energy and frame size. We soon show that this intuition is correct, in that all desired time average constraints are satisfied, and the time average power is within $O(1/V)$ of the optimal value $power^{opt}$. Hence, average power can be pushed arbitrarily close to optimal by using a sufficiently large value of V . The tradeoff is that average queue sizes grow with V , which impacts the convergence time required to satisfy the desired constraints.

The drift-plus-penalty ratio method was first developed for the context of restless bandit systems in [8][9]. The method was used for optimization of renewal systems in [1][2], which treat problems similar to those considered

in this paper. In the special case when all frame sizes are fixed and equal to one unit of time (a *time slot*), and when $V = 0$, the method reduces to observing queues $\mathbf{Q}[k]$ every slot k and taking actions to minimize a bound on $\mathbb{E}[\Delta[k]|\mathbf{Q}[k]]$. This is the rule that generates the classic max-weight scheduling algorithms for queue stability (without performance optimization), developed by Tassiulas and Ephremides in [10][11]. For systems with unit time slots but with $V > 0$, the *drift-plus-penalty ratio* technique reduces to the *drift-plus-penalty* technique of [12][13][14], which treats joint queue stability and penalty minimization in systems with unit size slots.

1) *Bounding the Drift-Plus-Penalty Ratio*: To construct an explicit algorithm, we first bound the drift-plus-penalty ratio.

Lemma 3: For all frames $k \in \{0, 1, 2, \dots\}$, all possible $\mathbf{Q}[k]$, and under any decisions for $(c[k], m[k], I[k])$, we have:

$$\begin{aligned} \frac{\mathbb{E}[\Delta[k] + V e[k]|\mathbf{Q}[k]]}{\mathbb{E}[D[k] + I[k]|\mathbf{Q}[k]]} &\leq \frac{B}{\mathbb{E}[D[k] + I[k]|\mathbf{Q}[k]]} + \frac{\mathbb{E}[V \hat{e}(c[k], m[k])|\mathbf{Q}[k]]}{\mathbb{E}[\hat{D}(c[k], m[k]) + I[k]|\mathbf{Q}[k]]} \\ &+ \frac{\sum_{n=1}^N Q_n[k] \mathbb{E}[\lambda_n(\hat{D}(c[k], m[k]) + I[k]) - 1_n[k]|\mathbf{Q}[k]]}{\mathbb{E}[\hat{D}(c[k], m[k]) + I[k]|\mathbf{Q}[k]]} \end{aligned} \quad (22)$$

where B is a constant that satisfies the following for all possible $\mathbf{Q}[k]$ and all policies:

$$B \geq \frac{1}{2} \sum_{n=1}^N \mathbb{E}[(\lambda_n(D[k] + I[k]) - 1_n[k])^2|\mathbf{Q}[k]]$$

Such a constant B exists by the second moment boundedness assumptions (1).

Proof: Note by iterated expectations that (similar to (7)):²

$$\mathbb{E}[D[k]|\mathbf{Q}[k]] = \mathbb{E}[\hat{D}(c[k], m[k])|\mathbf{Q}[k]] , \quad \mathbb{E}[e[k]|\mathbf{Q}[k]] = \mathbb{E}[\hat{e}(c[k], m[k])|\mathbf{Q}[k]]$$

Thus, the denominator is common for all terms of inequality (22), and it suffices to prove:

$$\mathbb{E}[\Delta[k]|\mathbf{Q}[k]] \leq B + \sum_{n=1}^N Q_n[k] \mathbb{E}[\lambda_n(\hat{D}(c[k], m[k]) + I[k]) - 1_n[k]|\mathbf{Q}[k]] \quad (23)$$

To this end, by squaring (18) and noting that $\max[x, 0]^2 \leq x^2$, we have for each n :

$$\begin{aligned} \frac{1}{2} Q_n[k+1]^2 &\leq \frac{1}{2} (Q_n[k] + \lambda_n(D[k] + I[k]) - 1_n[k])^2 \\ &= \frac{1}{2} Q_n[k]^2 + \frac{1}{2} (\lambda_n(D[k] + I[k]) - 1_n[k])^2 + Q_n[k](\lambda_n(D[k] + I[k]) - 1_n[k]) \end{aligned}$$

Summing the above over $n \in \{1, \dots, N\}$ and using the definition of $\Delta[k]$ gives:

$$\Delta[k] \leq \frac{1}{2} \sum_{n=1}^N (\lambda_n(D[k] + I[k]) - 1_n[k])^2 + \sum_{n=1}^N Q_n[k](\lambda_n(D[k] + I[k]) - 1_n[k])$$

Taking conditional expectations given $\mathbf{Q}[k]$ and using the bound B proves (23). \square

2) *The Task Scheduling Algorithm*: Our algorithm takes actions every frame to minimize the last two terms on the right-hand-side of the drift-plus-penalty ratio bound (22). The only part of these terms that we have control over on frame k (given the observed $\mathbf{Q}[k]$) is given below:

$$\frac{\mathbb{E}[V \hat{e}(c[k], m[k]) - \sum_{n=1}^N Q_n[k] 1_n[k]|\mathbf{Q}[k]]}{\mathbb{E}[\hat{D}(c[k], m[k]) + I[k]|\mathbf{Q}[k]]}$$

Recall from Lemma 1 that minimizing the above ratio of expectations is accomplished over a deterministic choice of $(c[k], m[k], I[k])$. Thus, every frame k we perform the following:

²In more detail, by iterated expectations we have $\mathbb{E}[D[k]|\mathbf{Q}[k]] = \mathbb{E}[\mathbb{E}[D[k]|(c[k], m[k]), \mathbf{Q}[k]]|\mathbf{Q}[k]]$, and $\mathbb{E}[D[k]|(c[k], m[k]), \mathbf{Q}[k]] = \mathbb{E}[D[k]|(c[k], m[k])]$ because $D[k]$ is conditionally independent of the past given the current $(c[k], m[k])$ used.

- Observe queues $\mathbf{Q}[k] = (Q_1[k], \dots, Q_N[k])$. Then choose $c[k] \in \{1, \dots, N\}$, $m[k] \in \mathcal{M}$, and $I[k]$ such that $0 \leq I[k] \leq I_{max}$ to minimize:

$$\frac{V\hat{e}(c[k], m[k]) - Q_{c[k]}[k]}{\hat{D}(c[k], m[k]) + I[k]} \quad (24)$$

- Update queues $Q_n[k]$ for each $n \in \{1, \dots, N\}$ via (18), using the $D[k]$, $I[k]$, and $1_n[k]$ values that result from the decisions $c[k]$, $m[k]$, $I[k]$ that minimized (24)

3) *Steps to minimize (24):* Here we elaborate on how to perform the minimization in (24) for each frame k . For each $c \in \{1, \dots, N\}$ and $m \in \mathcal{M}$, define $idle(c, m)$ as the value of $I[k]$ that minimizes (24), given that we have $(c[k], m[k]) = (c, m)$. It is easy to see that:

$$idle(c, m) = \begin{cases} 0 & \text{if } V\hat{e}(c, m) - Q_c[k] \leq 0 \\ I_{max} & \text{otherwise} \end{cases}$$

Now define $val(c, m)$ by:

$$val(c, m) = \frac{V\hat{e}(c, m) - Q_c[k]}{\hat{D}(c, m) + idle(c, m)}$$

Then we choose $(c[k], m[k])$ as the minimizer of $val(c, m)$ over $c \in \{1, \dots, N\}$ and $m \in \mathcal{M}$, breaking ties arbitrarily, and choose $I[k] = idle(c[k], m[k])$. Note that this algorithm chooses $I[k] = 0$ or $I[k] = I_{max}$ on every frame k . Nevertheless, it results in a *frame average* \bar{I} that approaches optimality for large V .

H. Performance of the Task Scheduling Algorithm

For simplicity, the performance theorem is presented in terms of zero initial conditions. It is assumed throughout that the problem (3)-(6) is feasible, so that it is possible to satisfy the constraints.

Theorem 1: Suppose $Q_n[0] = 0$ for all $n \in \{1, \dots, N\}$, and that the problem (3)-(6) is feasible. Then under the above task scheduling algorithm:

- a) For all frames $K \in \{1, 2, 3, \dots\}$ we have:³

$$\frac{\bar{e}[K]}{\bar{D}[K] + \bar{I}[K]} \leq power^{opt} + \frac{B}{V(\bar{D}[K] + \bar{I}[K])} \quad (25)$$

where B is defined in Lemma 3, $power^{opt}$ is the minimum power solution for the problem (3)-(6), and $\bar{e}[K]$, $\bar{D}[K]$, $\bar{I}[K]$ are defined by:

$$\bar{e}[K] \triangleq \frac{1}{K} \sum_{k=0}^{K-1} \mathbb{E}[e[k]] \quad , \quad \bar{D}[K] \triangleq \frac{1}{K} \sum_{k=0}^{K-1} \mathbb{E}[D[k]] \quad , \quad \bar{I}[K] \triangleq \frac{1}{K} \sum_{k=0}^{K-1} \mathbb{E}[I[k]]$$

- b) The desired constraints (20) and (21) are satisfied for all $n \in \{1, \dots, N\}$. Further, we have for each frame $K \in \{1, 2, 3, \dots\}$:

$$\frac{\mathbb{E}[|\mathbf{Q}[k]|]}{K} \leq \sqrt{\frac{2(B + V\beta)}{K}} \quad (26)$$

where $\|\mathbf{Q}[k]\| \triangleq \sqrt{\sum_{n=1}^N Q_n[k]^2}$ is the norm of the queue vector (being at least as large as each component $Q_n[k]$), and β is a constant that satisfies the following for all frames k :

$$\beta \geq \mathbb{E}[power^{opt}(D[k] + I[k]) - e[k]]$$

Such a constant β exists because the second moments (and hence first moments) are bounded.

In the special case of a deterministic system, all expectations of the above theorem can be removed, and the results hold deterministically for all frames K . Theorem 1 indicates that average power can be pushed arbitrarily close to $power^{opt}$, using the V parameter that affects an $O(1/V)$ performance gap given in (25). The tradeoff is that V increases the expected size of $\mathbb{E}[Q_n[K]]/K$ as shown in (26), which bounds the expected deviation from the n th constraint during the first K frames (recall (19) from Lemma 2). Under a mild additional “Slater-type” assumption that ensures all constraints can be satisfied with “ ϵ -slackness,” a stronger result on the virtual queues

³The right-hand-side in (25) can be simplified to $power^{opt} + B/(VD_{min})$, since all frames are at least D_{min} in size.

can be shown, namely, that the same algorithm yields queues with average size $O(V)$ [1]. This typically ensures a tighter constraint tradeoff than that given in (26). A related improved tradeoff is explored in more detail in Section III-C.

Proof: (Theorem 1 part (a)) Given $\mathbf{Q}[k]$ for frame k , our control decisions minimize the last two terms in the right-hand-side of the drift-plus-penalty ratio bound (22), and hence:

$$\begin{aligned} \frac{\mathbb{E}[\Delta[k] + Ve[k]|\mathbf{Q}[k]]}{\mathbb{E}[D[k] + I[k]|\mathbf{Q}[k]]} &\leq \frac{B}{\mathbb{E}[D[k] + I[k]|\mathbf{Q}[k]]} + \frac{\mathbb{E}[V\hat{e}(c^*[k], m^*[k])|\mathbf{Q}[k]]}{\mathbb{E}[\hat{D}(c^*[k], m^*[k]) + I^*[k]|\mathbf{Q}[k]]} \\ &+ \frac{\sum_{n=1}^N Q_n[k] \mathbb{E}[\lambda_n(\hat{D}(c^*[k], m^*[k]) + I^*[k]) - 1_n^*[k]|\mathbf{Q}[k]]}{\mathbb{E}[\hat{D}(c^*[k], m^*[k]) + I^*[k]|\mathbf{Q}[k]]} \end{aligned} \quad (27)$$

where $c^*[k]$, $m^*[k]$, $I^*[k]$, $1_n^*[k]$ are from any alternative (possibly randomized) decisions that can be made on frame k . Now recall the existence of stationary and randomized decisions that yield (15)-(16). In particular, these decisions are independent of queue backlog $\mathbf{Q}[k]$ and thus yield (from (15)):

$$\frac{\mathbb{E}[\hat{e}(c^*[k], m^*[k])|\mathbf{Q}[k]]}{\mathbb{E}[\hat{D}(c^*[k], m^*[k]) + I^*[k]|\mathbf{Q}[k]]} = \frac{\mathbb{E}[\hat{e}(c^*[k], m^*[k])]}{\mathbb{E}[\hat{D}(c^*[k], m^*[k]) + I^*[k]]} = power^{opt}$$

and for all $n \in \{1, \dots, N\}$ we have (from (16)):

$$\mathbb{E}[\lambda_n(\hat{D}(c^*[k], m^*[k]) + I^*[k]) - 1_n^*[k]|\mathbf{Q}[k]] = \mathbb{E}[\lambda_n(\hat{D}(c^*[k], m^*[k]) + I^*[k]) - 1_n^*[k]] \leq 0$$

Plugging the above into the right-hand-side of (27) yields:

$$\frac{\mathbb{E}[\Delta[k] + Ve[k]|\mathbf{Q}[k]]}{\mathbb{E}[D[k] + I[k]|\mathbf{Q}[k]]} \leq \frac{B}{\mathbb{E}[D[k] + I[k]|\mathbf{Q}[k]]} + Vpower^{opt}$$

Rearranging terms gives:

$$\mathbb{E}[\Delta[k] + Ve[k]|\mathbf{Q}[k]] \leq B + Vpower^{opt}\mathbb{E}[D[k] + I[k]|\mathbf{Q}[k]]$$

Taking expectations of the above (with respect to the random $\mathbf{Q}[k]$) and using the law of iterated expectations gives:

$$\mathbb{E}[\Delta[k] + Ve[k]] \leq B + Vpower^{opt}\mathbb{E}[D[k] + I[k]] \quad (28)$$

The above holds for all $k \in \{0, 1, 2, \dots\}$. Fixing a positive integer K and summing (28) over $k \in \{0, 1, \dots, K-1\}$ yields, by the definition $\Delta[k] = L[k+1] - L[k]$:

$$\mathbb{E}[L[K] - L[0]] + V \sum_{k=0}^{K-1} \mathbb{E}[e[k]] \leq BK + Vpower^{opt} \sum_{k=0}^{K-1} \mathbb{E}[D[k] + I[k]]$$

Noting that $L[0] = 0$ and $L[K] \geq 0$ and using the definitions of $\bar{e}[K]$, $\bar{D}[K]$, $\bar{I}[K]$ yields:

$$VK\bar{e}[K] \leq BK + VKpower^{opt}(\bar{D}[K] + \bar{I}[K])$$

Rearranging terms yields the result of part (a). □

Proof: (Theorem 1 part (b)) To prove part (b), note from (28) we have:

$$\mathbb{E}[\Delta[k]] \leq B + V\mathbb{E}[power^{opt}(D[k] + I[k]) - e[k]] \leq B + V\beta$$

Summing the above over $k \in \{0, 1, \dots, K-1\}$ gives:

$$\mathbb{E}[L[K]] - \mathbb{E}[L[0]] \leq (B + V\beta)K$$

Using the definition of $L[K]$ and noting that $L[0] = 0$ gives:

$$\sum_{l=1}^K \mathbb{E}[Q_n[l]^2] \leq 2(B + V\beta)K$$

Thus, we have $\mathbb{E} [||\mathbf{Q}[K]||^2] \leq 2(B+V\beta)K$. Jensen's inequality for $f(x) = x^2$ ensures $\mathbb{E} [||\mathbf{Q}[k]||]^2 \leq \mathbb{E} [||\mathbf{Q}[k]||^2]$, and so for all positive integers K :

$$\mathbb{E} [||\mathbf{Q}[k]||]^2 \leq 2(B + V\beta)K \quad (29)$$

Taking a square root of both sides of (29) and dividing by K proves (26). From (26) we have for each $n \in \{1, \dots, N\}$:

$$\lim_{K \rightarrow \infty} \frac{\mathbb{E} [Q_n[k]]}{K} \leq \lim_{K \rightarrow \infty} \frac{\mathbb{E} [||\mathbf{Q}[K]||]}{K} \leq \lim_{K \rightarrow \infty} \frac{\sqrt{2(B + V\beta)}}{\sqrt{K}} = 0$$

and hence by Lemma 2 we know constraint (21) holds. Further, in [15] it is shown that (29) together with the fact that second moments of queue changes are bounded implies $\lim_{k \rightarrow \infty} Q_n[k]/k = 0$ with probability 1. Thus, (20) holds. \square

I. Simulation

We first simulate the task scheduling algorithm for the simple deterministic system with one class and one constraint, as described in Section II-D. The $\hat{e}(m)$ and $\hat{D}(m)$ functions are defined in (8)-(9), and the goal is to minimize average power subject to a processing rate constraint $1/(\bar{D} + \bar{I}) \geq 0.2$. We already know the optimal power is $power^{opt} = 7/15 \approx 0.466667$. We expect the algorithm to approach this optimal power as V is increased, and to approach the desired behavior of using $I[k] = 0$ for all k , meeting the constraint with equality, and using $m[k] = 1$ for $1/3$ of the frames. This is indeed what happens, although in this simple case the algorithm seems insensitive to the V parameter and locks into a desirable periodic schedule even for very low (but positive) V values. Using $V = 1$ and one million frames, the algorithm gets average power 0.466661, uses $m[k] = 1$ a fraction of time 0.333340, has average idle time $\bar{I} = 0.000010$, and yields a processing rate 0.199999 (almost exactly equal to the desired constraint of 0.2). Increasing V yields similar performance. The constraint is still satisfied when we decrease the value of V , but average power degrades (being 0.526316 for $V = 0$).

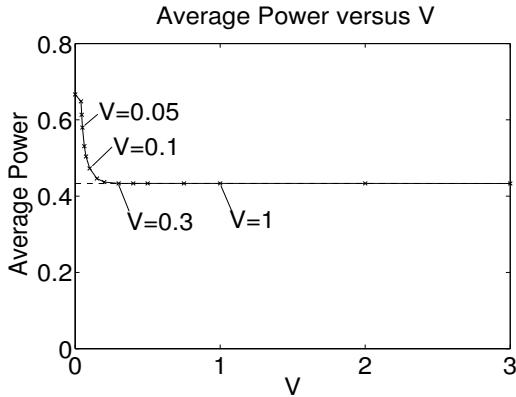


Fig. 3. Average power versus V .

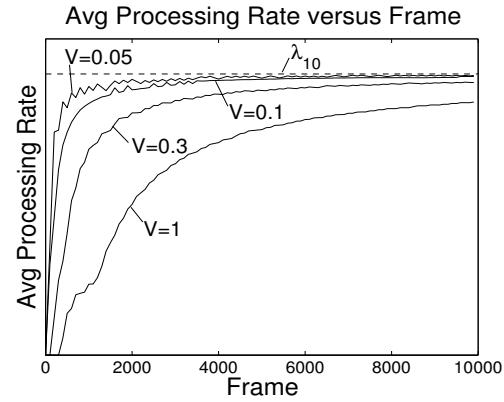


Fig. 4. Processing Rate \bar{I}_{10}/\bar{T} versus frame index.

We next consider a system with 10 classes of tasks and two processing modes. The energy and delay characteristics for each class $i \in \{1, \dots, 10\}$ and mode $m \in \{1, 2\}$ are:

$$\text{Mode 1: } (\hat{e}(i, 1), \hat{D}(i, 1)) = (1i, 5i) \quad (30)$$

$$\text{Mode 2: } (\hat{e}(i, 2), \hat{D}(i, 2)) = (2i, 3i) \quad (31)$$

so that mode 1 uses less energy but takes longer than mode 2, and the computational requirements for each class increase with $i \in \{1, \dots, 10\}$. We assume desired rates are given by $\lambda_i = \rho/(30i)$ for $i \in \{1, \dots, 10\}$, for some positive value ρ . The problem is feasible whenever $\rho \leq 1$. We use $\rho = 0.8$ and run the simulation for 10 million frames. Fig. 3 shows the resulting average power as V is varied between 0 and 3, which converges to near optimal after $V = 0.3$. All 10 processing rate constraints are met within 5 decimal points of accuracy after the 10 million frames. An illustration of how convergence time is affected by the V parameter is shown in Fig. 4, which illustrates the average processing rate \bar{I}_{10}/\bar{T} for class 10, as compared to the desired constraint $\lambda_{10} = \rho/300$. It is seen, for example, that convergence is faster for $V = 0.05$ than for $V = 1$. Convergence times can be improved using non-zero initial queue backlog and the theory of *place holder backlog* in [1], although we omit this topic for brevity.

III. OPTIMIZATION WITH GENERAL ATTRIBUTES

This section generalizes the problem to allow time average optimization for abstract *attributes*. Consider again a frame-based system with frame index $k \in \{0, 1, 2, \dots\}$. Every frame k , the controller makes a *control action* $\alpha[k]$, chosen within an abstract set \mathcal{A} of allowable actions. The action $\alpha[k]$ affects the *frame size* $T[k]$ and an *attribute vector* $\mathbf{y}[k] = (y_0[k], y_1[k], \dots, y_L[k])$. Specifically, assume these are random functions that are conditionally independent of the past given the current $\alpha[k]$ decision, with mean values given by functions $\hat{T}(\alpha)$ and $\hat{y}_l(\alpha)$ for all $\alpha \in \mathcal{A}$:

$$\hat{T}(\alpha) = \mathbb{E}[T[k]|\alpha[k] = \alpha], \quad \hat{y}_l(\alpha) = \mathbb{E}[y_l[k]|\alpha[k] = \alpha]$$

Similar to the previous section, it is assumed there is a minimum frame size $T_{min} > 0$ such that $T[k] \geq T_{min}$ for all k , and that second moments are bounded by a constant σ^2 , regardless of the policy $\alpha[k]$. The joint distribution of $(T[k], y_0[k], y_1[k], \dots, y_L[k])$ is otherwise arbitrary.

Define frame averages \bar{T} and \bar{y}_l by:

$$\bar{T} = \lim_{K \rightarrow \infty} \frac{1}{K} \sum_{k=0}^{K-1} T[k], \quad \bar{y}_l = \lim_{K \rightarrow \infty} \frac{1}{K} \sum_{k=0}^{K-1} y_l[k]$$

As discussed in Section II-B, the value \bar{y}_l/\bar{T} represents the *time average* associated with attribute $y_l[k]$. The general problem is then:

$$\text{Minimize: } \bar{y}_0/\bar{T} \quad (32)$$

$$\text{Subject to: } \bar{y}_l/\bar{T} \leq c_l \quad \forall l \in \{1, \dots, L\} \quad (33)$$

$$\alpha[k] \in \mathcal{A} \quad \forall k \in \{0, 1, 2, \dots\} \quad (34)$$

where c_1, \dots, c_L are given constants that specify the desired time average constraints.

A. Mapping to the Task Scheduling Problem

To illustrate the generality of this framework, this subsection uses the new notation to exactly represent the task scheduling problem from Section II. For that problem, one can define the control action $\alpha[k]$ to have the form $\alpha[k] = (c[k], m[k], I[k])$, and the action space \mathcal{A} is then the set of all (c, m, I) such that $c \in \{1, \dots, N\}$, $m \in \mathcal{M}$, and $0 \leq I \leq I_{max}$.

The frame size is $T[k] = D[k] + I[k]$, and $\hat{T}(\alpha[k])$ is given by:

$$\hat{T}(\alpha[k]) = \hat{D}(c[k], m[k]) + I[k]$$

We then define $y_0[k]$ as the energy expended in frame k , so that $y_0[k] = e[k]$ and $\hat{y}_0(\alpha[k]) = \hat{e}(c[k], m[k])$. There are N constraints, so define $L = N$. To express the desired constraints $\bar{y}_n/\bar{T} \geq \lambda_n$ in the form $\bar{y}_n/\bar{T} \leq c_n$, one can define $y_n[k] = -1_n[k]$ and $c_n = -\lambda_n$ for each $n \in \{1, \dots, N\}$, and $\hat{y}_n(\alpha[k]) = -1_n[k]$. Alternatively, one could define $y_n[k] = \lambda_n T[k] - 1_n[k]$ and enforce the constraint $\bar{y}_n \leq 0$ for all $n \in \{1, \dots, N\}$.

This general setup provides more flexibility. For example, suppose the idle state does not use 0 energy, but operates at a low power p_{idle} and expends total energy $p_{idle}I[k]$ on frame k . Then total energy for frame k can be defined as $y_0[k] = e[k] + p_{idle}I[k]$, where $e[k]$ is the energy spent in the busy period. The setup can also handle systems with multiple idle mode options, each providing a different energy savings but incurring a different wakeup time.

B. The General Algorithm

The algorithm for solving the general problem (32)-(34) is described below. Each constraint $\bar{y}_l \leq c_l \bar{T}$ in (33) is treated using a virtual queue $Q_l[k]$ with update:

$$Q_l[k+1] = \max[Q_l[k] + y_l[k] - c_l T[k], 0] \quad \forall l \in \{1, \dots, L\} \quad (35)$$

Defining $L[k]$ and $\Delta[k]$ as before (in Section II-G) leads to the following bound on the drift-plus-penalty ratio, which can be proven in a manner similar to Lemma 3:

$$\frac{\mathbb{E}[\Delta[k] + V y_0[k] | \mathbf{Q}[k]]}{\mathbb{E}[T[k] | \mathbf{Q}[k]]} \leq \frac{B}{\mathbb{E}[T[k] | \mathbf{Q}[k]]} + \frac{\mathbb{E}\left[V \hat{y}_0(\alpha[k]) + \sum_{l=1}^L Q_l[k] \hat{y}_l(\alpha[k]) | \mathbf{Q}[k]\right]}{\mathbb{E}[\hat{T}(\alpha[k]) | \mathbf{Q}[k]]} \quad (36)$$

where B is a constant that satisfies the following for all $\mathbf{Q}[k]$ and all possible actions $\alpha[k]$:

$$B \geq \frac{1}{2} \sum_{l=1}^L \mathbb{E}[(y_l[k] - c_l T[k])^2 | \mathbf{Q}[k]]$$

Every frame, the controller observes queues $\mathbf{Q}[k]$ and takes an action $\alpha[k] \in \mathcal{A}$ that minimizes the second term on the right-hand-side of (36). We know from Lemma 1 that minimizing the ratio of expectations is accomplished by a deterministic selection of $\alpha[k] \in \mathcal{A}$. The resulting algorithm is:

- Observe $\mathbf{Q}[k]$ and choose $\alpha[k] \in \mathcal{A}$ to minimize (breaking ties arbitrarily):

$$\frac{V \hat{y}_0(\alpha[k]) + \sum_{l=1}^L Q_l[k] \hat{y}_l(\alpha[k])}{\hat{T}(\alpha[k])} \quad (37)$$

- Update virtual queues $Q_l[k]$ for each $l \in \{1, \dots, L\}$ via (35).

One subtlety is that the expression (37) may not have an achievable minimum over the general (possibly infinite) set \mathcal{A} (for example, the infimum of the function $f(x) = x$ over the open interval $0 < x < 1$ is not achievable over that interval). This is no problem: Our algorithm in fact works for any *approximate* minimum that is an additive constant C away from the exact infimum every frame k (for any arbitrarily large constant $C \geq 0$). This effectively changes the “ B ” constant in our performance bounds to a new constant “ $B + C$ ” [1]. Let $\text{ratio}^{\text{opt}}$ represent the optimal ratio of \bar{y}_0/\bar{T} for the problem (32)-(34). As before, it can be shown that if the problem is feasible (so that there exists an algorithm that can achieve the constraints (33)-(34)), then any C -additive approximation of the above algorithm satisfies all desired constraints and yields $\bar{y}_0/\bar{T} \leq \text{ratio}^{\text{opt}} + O(1/V)$, which can be pushed arbitrarily close to $\text{ratio}^{\text{opt}}$ as V is increased, with the same tradeoff in the queue sizes (and hence convergence times) with V . The proof of this is similar to that of Theorem 1, and is omitted for brevity (see [1][2] for the full proof).

C. Random Task Arrivals and Flow Control

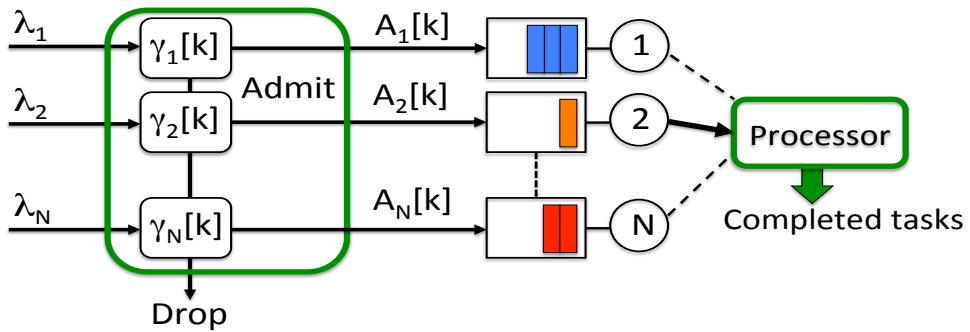


Fig. 5. A task processing system with random arrivals and flow control.

Again consider a system with N classes of tasks, as in Section II. Each frame k again has a busy period of duration $D[k]$ and an idle period of duration $I[k]$ as in Fig. 1. However, rather than always having tasks available for processing, this subsection assumes tasks arrive *randomly* with rates $(\lambda_1, \dots, \lambda_N)$, where λ_n is the rate of task arrivals per unit time (see Fig. 5). At the beginning of each busy period, the controller chooses a variable $c[k]$ that specifies which type of task is performed. However, $c[k]$ can now take values in the set $\{0, 1, \dots, N\}$, where $c[k] = 0$ is a *null* choice that selects no task on frame k . If $c[k] = 0$, the busy period has some positive size D_0 and may spend a small amount of energy to power the electronics, but does not process any task. The mode selection

variable $m[k]$ takes values in the same set \mathcal{M} as before. The idle time variable $I[k]$ is again chosen in the interval $[0, I_{max}]$.

Further, for each $n \in \{1, 2, \dots, N\}$ we introduce *flow control variables* $\gamma_n[k]$, chosen in the interval $0 \leq \gamma_n[k] \leq 1$. The variable $\gamma_n[k]$ represents the probability of admitting each new randomly arriving task of class n on frame k (see Fig. 5). This enables the system to drop tasks if the raw arrival rates $(\lambda_1, \dots, \lambda_N)$ cannot be supported. Let $\gamma[k] = (\gamma_1[k], \dots, \gamma_N[k])$ be the vector of these variables.

We thus have $\alpha[k] = (c[k], m[k], I[k], \gamma[k])$, with action space \mathcal{A} being the set of all (c, m, I, γ) such that $c \in \{0, 1, \dots, N\}$, $m \in \mathcal{M}$, $0 \leq I \leq I_{max}$, and $0 \leq \gamma_n \leq 1$ for $n \in \{1, \dots, N\}$. Define $e[k]$ and $D[k]$ as the energy and busy period duration for frame k . Assume $e[k]$ depends only on $(c[k], m[k], I[k])$, and $D[k]$ depends only on $(c[k], m[k])$, with averages given by functions $\hat{e}(c[k], m[k], I[k])$ and $\hat{D}(c[k], m[k])$:

$$\hat{e}(c[k], m[k], I[k]) = \mathbb{E}[e[k]|c[k], m[k], I[k]] \quad , \quad \hat{D}(c[k], m[k]) = \mathbb{E}[D[k]|c[k], m[k]]$$

Finally, for each $n \in \{1, \dots, N\}$, define $A_n[k]$ as the random number of new arrivals admitted on frame k , which depends on the total frame size $D[k] + I[k]$ and the admission probability $\gamma_n[k]$. Formally, assume the arrival vector $(A_1[k], \dots, A_N[k])$ is conditionally independent of the past given the current $\alpha[k]$ used, with expectations:

$$\mathbb{E}[A_n[k]|\alpha[k]] = \lambda_n \gamma_n[k] [\hat{D}(c[k], m[k]) + I[k]] \quad (38)$$

The assumption on independence of the past holds whenever arrivals are independent and Poisson, or when all frame sizes are an integer number of fixed size slots, and arrivals are independent and identically distributed (i.i.d.) over slots with some general distribution.

We seek to maximize a weighted sum of admission rates subject to supporting all of the admitted tasks, and to maintaining average power to within a given positive constant P_{av} :

$$\text{Maximize: } \frac{\sum_{n=1}^N w_n \bar{A}_n}{\bar{D} + \bar{I}} \quad (39)$$

$$\text{Subject to: } \bar{A}_n / (\bar{D} + \bar{I}) \leq \bar{1}_n / (\bar{D} + \bar{I}) \quad \forall n \in \{1, \dots, N\} \quad (40)$$

$$\frac{\bar{e}}{\bar{D} + \bar{I}} \leq P_{av} \quad (41)$$

$$\alpha[k] \in \mathcal{A} \quad \forall k \in \{0, 1, 2, \dots\} \quad (42)$$

where (w_1, \dots, w_N) are a collection of positive weights that prioritize the different classes in the optimization objective.

We thus have $L = N + 1$ constraints. To treat this problem, define $T[k] = D[k] + I[k]$, so that $\hat{T}(\alpha[k]) = \hat{D}(c[k], m[k]) + I[k]$. Further define $y_0[k]$, $y_n[k]$ for $n \in \{1, \dots, N\}$, and $x[k]$ as:

$$\begin{aligned} y_0[k] = - \sum_{n=1}^N w_n A_n[k] &\implies \hat{y}_0(\alpha[k]) = - \sum_{n=1}^N w_n \lambda_n \gamma_n[k] [\hat{D}(c[k], m[k]) + I[k]] \\ y_n[k] = A_n[k] - 1_n[k] &\implies \hat{y}_n(\alpha[k]) = \lambda_n \gamma_n[k] [\hat{D}(c[k], m[k]) + I[k]] - 1_n[k] \\ x[k] = e[k] - [D[k] + I[k]] P_{av} &\implies \hat{x}(\alpha[k]) = \hat{e}(c[k], m[k], I[k]) - [\hat{D}(c[k], m[k]) + I[k]] P_{av} \end{aligned}$$

Then:

- Minimizing \bar{y}_0 / \bar{T} is equivalent to (39).
- The constraints $\bar{y}_n \leq 0$ for $n \in \{1, \dots, N\}$ are equivalent to (40).
- The constraint $\bar{x} \leq 0$ is equivalent to (41).

Note that the above problem does not specify any explicit queueing for the randomly arriving tasks. The algorithm will in fact construct explicit queues (so that the virtual queues can be viewed as actual queues). Note also that the constraint $\alpha[k] \in \mathcal{A}$ does not allow restrictions on actions based on the queue state, such as when the queue is empty or not. Thus, in principle, we allow the possibility of “processing” a task of class n even when there is no such task available. In this case, we assume this processing is still costly, in that it incurs time equal to $\hat{D}(c[k], m[k])$ and energy equal to $\hat{e}(c[k], m[k], I[k])$. Our algorithm will naturally learn to avoid the inefficiencies associated with such actions.

1) *The Dynamic Algorithm for Random Task Arrivals:* To enforce the constraints $\bar{y}_n \leq 0$ for each $n \in \{1, \dots, N\}$, define queue $Q_n[k]$ with update:

$$Q_n[k+1] = \max[Q_n[k] + A_n[k] - 1_n[k], 0] \quad (43)$$

To enforce $\bar{x} \leq 0$, define a virtual queue $Z[k]$ with update:

$$Z[k+1] = \max[Z[k] + e[k] - [D[k] + I[k]]P_{av}, 0] \quad (44)$$

It can be seen that the queue update (43) is the same as that of an *actual queue* for class n tasks, with random task arrivals $A_n[k]$ and task service $1_n[k]$. The minimization of (37) then becomes the following: Every frame k , observe queues $Q[k]$ and $Z[k]$. Then choose $c[k] \in \{0, 1, \dots, N\}$, $m[k] \in \mathcal{M}$, $I[k] \in [0, I_{max}]$, and $\gamma_n[k] \in [0, 1]$ for all $n \in \{1, \dots, N\}$ to minimize:

$$\begin{aligned} & \frac{-V \sum_{n=1}^N w_n \lambda_n \gamma_n[k] [\hat{D}(c[k], m[k]) + I[k]]}{\hat{D}(c[k], m[k]) + I[k]} \\ & + \frac{\sum_{n=1}^N Q_n[k] (\lambda_n \gamma_n[k] [\hat{D}(c[k], m[k]) + I[k]] - 1_n[k])}{\hat{D}(c[k], m[k]) + I[k]} \\ & + \frac{Z[k] (\hat{e}(c[k], m[k], I[k]) - [\hat{D}(c[k], m[k]) + I[k]] P_{av})}{\hat{D}(c[k], m[k]) + I[k]} \end{aligned} \quad (45)$$

After a simplifying cancellation of terms, it is easy to see that the $\gamma_n[k]$ decisions can be separated from all other decisions (see Exercise 2). The resulting algorithm then observes queues $Q[k]$ and $Z[k]$ every frame k and performs the following:

- (Flow Control) For each $n \in \{1, \dots, N\}$, choose $\gamma_n[k]$ as:

$$\gamma_n[k] = \begin{cases} 1 & \text{if } Q_n[k] \leq V w_n \\ 0 & \text{otherwise} \end{cases} \quad (46)$$

- (Task Scheduling) Choose $c[k] \in \{0, 1, \dots, N\}$, $m[k] \in \mathcal{M}$, $I[k] \in [0, I_{max}]$ to minimize:

$$\frac{Z[k] \hat{e}(c[k], m[k], I[k]) - \sum_{n=1}^N Q_n[k] 1_n[k]}{\hat{D}(c[k], m[k]) + I[k]} \quad (47)$$

- (Queue Update) Update $Q_n[k]$ for each $n \in \{1, \dots, N\}$ by (43), and update $Z[k]$ by (44).

The minimization problem (47) is similar to (24), and can be carried out in the same manner as discussed in Section II-G. A key observation about the above algorithm is that *it does not require knowledge of the arrival rates* $(\lambda_1, \dots, \lambda_N)$. Indeed, the λ_n terms cancel out of the minimization, so that the flow control variables $\gamma_n[k]$ in (46) make “bang-bang” decisions that admit all newly arriving tasks of class n on frame k if $Q_n[k] \leq V w_n$, and admit none otherwise. This property makes the algorithm naturally adaptive to situations when the arrival rates change, as shown in the simulations of Section III-D.

Note that if $\hat{e}(0, m, I) < \hat{e}(c, m, I)$ for all $c \in \{1, \dots, N\}$, $m \in \mathcal{M}$, $I \in [0, I_{max}]$, so that the energy associated with processing *no task* is less than the energy of processing any class $c \neq 0$, then the minimization in (47) will never select a class c such that $Q_c[k] = 0$. That is, the algorithm naturally will never select a class for which no task is available.

2) *Deterministic Queue Bounds and Constraint Violation Bounds:* In addition to satisfying the desired constraints and achieving a weighted sum of admitted rates that is within $O(1/V)$ of optimality, the flow control structure of the task scheduling algorithm admits *deterministic queue bounds*. Specifically, assume all frame sizes are bounded by a constant T_{max} , and that the raw number of class n arrivals per frame (before admission control) is at most $A_{n,max}$. By the flow control policy (46), new arrivals of class n are only admitted if $Q_n[k] \leq V w_n$. Thus, assuming that $Q_n[0] \leq V w_n + A_{n,max}$, we must have $Q_n[k] \leq V w_n + A_{n,max}$ for all frames $k \in \{0, 1, 2, \dots\}$. This specifies a worst-case queue backlog that is $O(V)$, which establishes an explicit $[O(1/V), O(V)]$ performance-backlog tradeoff that is superior to that given in (26).

With mild additional structure on the $\hat{e}(c, m, I)$ function, the deterministic bounds on queues $Q_n[k]$ lead to a deterministic bound on $Z[k]$, so that one can compute a value Z_{max} , of size $O(V)$, such that $Z[k] \leq Z_{max}$ for all k . This is explored in Exercise 3 (see also [14]).

D. Simulations and Adaptiveness of Random Task Scheduling

Here we simulate the dynamic task scheduling and flow control algorithm (46)-(47), using the 10-class system model defined in Section II-I with $\hat{e}(c, m)$ functions given in (30)-(31). For consistency with that model, we remove the $c[k] = 0$ option, so that the decision (31) chooses $c[k] = 1, m[k] = 1$, incurring one unit of energy, in case no tasks are available. Arrivals are from independent Bernoulli processes with rates $\lambda_i = \rho/(30i)$ for each class $i \in \{1, \dots, 10\}$, with $\rho = 0.8$. We use weights $w_n = 1$ for all n , so that the objective is to maximize total throughput, and $P_{av} = 0.5$, which we know is feasible from results in Fig. 3 of Section II-I. Thus, we expect the algorithm to learn to admit everything, so that the admission rate approaches the total arrival rate as V is increased. We simulate for 10 million frames, using V in the interval from 0 to 200. Results are shown in Figs. 6 and 7. Fig. 6 shows the algorithm learns to admit everything for large V (100 or above), and Fig. 7 plots the resulting average queue size (in number of tasks) per queue, together with the deterministic bound $Q_n[k] \leq V + 60$ (where $A_{n,max} = 60$ in this case because there is at most one arrival per slot, and the largest possible frame is $D_{max} + I_{max} = 50 + 10 = 60$). The average power constraint was satisfied (with slackness) for all cases. The average queue size in Fig. 7 grows linearly with V until $V = 100$, when it saturates by admitting everything. The saturation value is the average queue size associated with admitting the raw arrival rates directly.

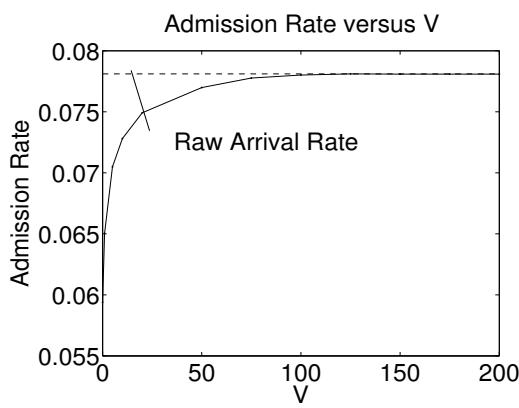


Fig. 6. Total admission rate versus V .

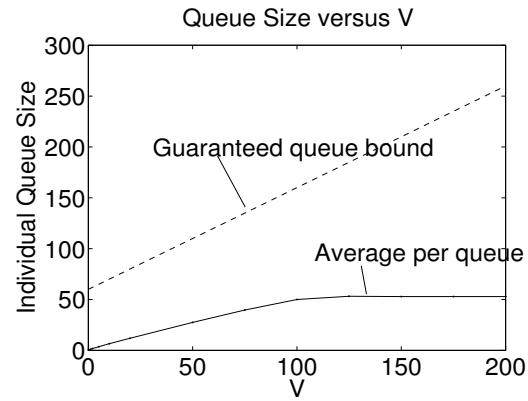


Fig. 7. Average queue size per queue, and the deterministic bound from Section III-C.2.

We now illustrate that the algorithm is robust to abrupt rate changes. We consider the same system as before, run over 10 million frames with $V = 100$. However, we break the simulation timeline into three equal size phases. During the first and third phase, we use arrival rates $\lambda_i = \rho_1/(30i)$ for $i \in \{1, \dots, 10\}$, where $\rho_1 = 0.8$ as in the previous experiment. Then, during the second (middle) phase, we double the rates to $\lambda_i = \rho_2/(30i)$, where $\rho_2 = 1.6$. Because $\rho_2 > 1$, these rates are *infeasible* and the algorithm must learn to optimally drop tasks so as to maximize the admission rate subject to the power constraint. Recall that the algorithm is unaware of the arrival rates and must adapt to the existing system conditions. The results are shown in Figs. 8 and 9. Fig. 8 shows a moving average admission rate versus time. During the first and third phases of the simulation, we have $\rho_1 = 0.8$ and the admitted rates are close to the raw arrival rate (shown as the lower dashed horizontal line). During the middle interval (with $\rho_2 = 1.6$), the algorithm quickly adapts to the increased arrivals and yields admitted rates that are close to those that should be used in a system with loading $\rho_2 = 1.6$ always (shown as the higher dashed horizontal line). Fig. 9 plots the corresponding moving average queue backlog per queue. The lower dashed horizontal line indicates the value of average backlog that would be achieved in a system with loading $\rho_1 = 0.8$ always. Also shown is the deterministic queue bound $V + 60 = 160$, which holds for all frames, regardless of the raw arrival rates.

E. Task Scheduling: Extensions and Further Reading

This section considered minimizing time averages subject to time average constraints. Extended techniques for minimizing convex functions of time average vectors (or maximizing concave functions) subject to similar constraints

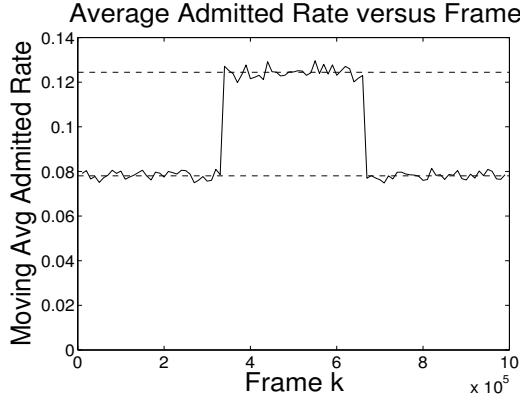


Fig. 8. Moving average admission rate versus frame index k for the system with abrupt rate changes.

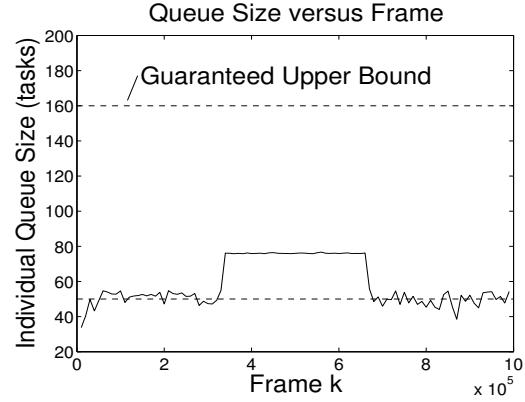


Fig. 9. Individual queue size (average and guaranteed bound) versus frame index k for the system with abrupt rate changes.

are treated in the renewal optimization theory of [1][2]. This is useful for extending the flow control problem (39)-(42) to optimize a sum of concave utility functions $\phi_n(x)$ of the time average admission rates for each class:

$$\sum_{n=1}^N \phi_n \left(\bar{A}_n / (\bar{D} + \bar{I}) \right)$$

Using logarithmic functions $\phi_n(x)$ leads to *proportional fairness* [16], while other fairness properties are achieved with other concave functions [17][18][19]. Related utility optimization for particular classes of wireless systems with fixed size frames of T slots, and with probabilistic reception every slot, is treated using a different technique in [20][21]. Related work on utility optimal scheduling for single-slot data networks is treated using convex optimization in [22][23][24][25], stochastic “primal-dual” algorithms in [26][27][28][29][30], and stochastic “dual” algorithms in [12][31][32][13][33].

For the problem with random arrivals, queueing delay can often be improved by changing the constraint $\bar{A}_n \leq \bar{I}_n$ of (40) to $\bar{A}_n + \epsilon \leq \bar{I}_n$, for some $\epsilon > 0$. This specifies that the output processing rate should be ϵ larger than the arrival rate. However, unlike the case $\epsilon = 0$, such a constraint yields a queue $Q_n[k]$ that contains some “fake” tasks, and can lead to decisions that serve the queue when no actual tasks are available. This can be overcome using the theory of ϵ -persistent service queues in [34][35].

Note that our flow control algorithm rewards admission of new data on each frame. This fits the general framework of this section, where attributes $y_l[k]$ are (possibly random) functions of the control action $\alpha[k] \in \mathcal{A}$. One might attempt to solve the problem by defining a reward upon *completion* of service. This does not fit our framework: That is because the algorithm might “complete” a service in a queue that is empty, simply to accumulate the “fake” reward. One could eliminate fake rewards by an augmented model that allows rewards $\hat{y}_l(\alpha[k], Q[k])$ and/or action spaces $\mathcal{A}(Q[k])$ that depend on the current backlog, although this creates a much more complex Markov decision model that is easily avoided by defining rewards at admission. However, there are some problems that require such a reward structure, such as problems of stock market trading where prior ownership of a particular stock is required to reap the benefits of selling at an observed desirable price. These problems can be treated with a modified Lyapunov function of the form $L[k] = \sum_{n=1}^N (Q_n[k] - \theta)^2$, which pushes backlog towards a non-empty state θ . This approach is used for stock trading [36], inventory control [37], energy harvesting [38], and smart-grid energy management [39]. The first algorithms for optimal energy harvesting used a related technique [40], and related problems in processing networks that assemble components by combining different sub-components are treated in [41][42].

F. Exercises for Section III

Exercise 1: Consider a system with N classes, each of which always has new tasks available. Every frame k we choose a class $c[k] \in \{1, \dots, N\}$, and select a single task of that class for processing. We can process using mode $m[k] \in \mathcal{M}$. The result yields frame size $\hat{T}(c[k], m[k])$, energy $\hat{e}(c[k], m[k])$, and processing quality

$\hat{q}(c[k], m[k])$. Design an algorithm that selects $c[k]$ and $m[k]$ every frame k to maximize \bar{q}/\bar{T} subject to an average power constraint $\bar{e}/\bar{T} \leq P_{av}$ and to a processing rate constraint of at least λ_n for each class n .

Exercise 2: Show that minimization of (45) leads to the separable flow control and task scheduling algorithm of (46)-(47).

Exercise 3: Consider the flow control and task scheduling algorithm (46)-(47), and recall that $Q_n[k] \leq Q_{max}$ for all $n \in \{1, \dots, N\}$ and all frames k , where $Q_{max} \triangleq \max_{n \in \{1, \dots, N\}} [Vw_n + A_{n,max}]$. Suppose there is a constant $e_{min} > 0$ such that $\hat{e}(c, m, I) \geq e_{min}$ for all $c \neq 0$, and that $\hat{e}(0, m, I) = 0$ for all $m \in \mathcal{M}$, $I \in [0, I_{max}]$.

- a) Show that the minimization of (47) chooses $c[k] = 0$ whenever $Z[k] > Q_{max}/e_{min}$.
- b) Suppose there is a constant e_{max} such that $e[k] \leq e_{max}$ for all k . Conclude that $Z[k] \leq Z_{max} \triangleq Q_{max}/e_{min} + e_{max}$ for all frames k .
- c) Use the queueing equation (44) to show that over any sequence of K frames $\{j, j+1, \dots, j+K-1\}$ (starting at some frame j), the total energy usage satisfies:

$$\sum_{k=j}^{j+K-1} e[k] \leq P_{av} \sum_{k=j}^{j+K-1} T[k] + Z_{max}$$

where $T[k]$ is the size of frame k . Hint: Make an argument similar to the proof of Lemma 2.

Exercise 4: Consider the same general attributes $y_l[k]$ of Section III, with $\hat{y}_l(\alpha[k])$ for $\alpha[k] \in \mathcal{A}$. State the algorithm for solving the problem of minimizing \bar{y}_0 subject to $\bar{y}_l/\bar{T} \leq c_l$. Hint: Define appropriate attributes $x_l[k]$ on a system with an “effective” frame size of 1 every frame, and minimize $\bar{x}_0/1$ subject to $\bar{x}_l/1 \leq 0$ for $l \in \{1, \dots, L\}$.

Exercise 5: Modify the task scheduling algorithm of Section II-G to allow the controller to serve more than one task per frame. Specifically, every frame k the controller chooses a *service action* $s[k]$ from a set \mathcal{S} of possible actions. Each service action $s \in \mathcal{S}$ determines a *clearance vector* $\mathbf{c}(s) = (c_1(s), \dots, c_N(s))$, where $c_i(s)$ is the number of tasks of type i served if action s is used on a given frame. It also incurs a delay $\hat{D}(s)$ and energy $\hat{e}(s)$.

Exercise 6: Consider the linear fractional problem of finding a vector (x_1, \dots, x_M) to solve:

$$\begin{aligned} \text{Minimize: } & \frac{a_0 + \sum_{i=1}^M a_i x_i}{b_0 + \sum_{i=1}^M b_i x_i} \\ \text{Subject to: } & \sum_{i=1}^M c_{il} x_i \leq d_l \quad \forall l \in \{1, \dots, L\} \\ & 0 \leq x_i \leq 1 \quad \forall i \in \{1, \dots, M\} \end{aligned}$$

Assume constants a_i, b_i, c_{il}, d_l are given, that $b_0 > 0$, and that $b_i \geq 0$ for all $i \in \{1, \dots, M\}$. Treat this as a time average problem with action $\alpha[k] = (x_1[k], \dots, x_M[k])$, action space $\mathcal{A} = \{(x_1, \dots, x_M) | 0 \leq x_i \leq 1 \ \forall i \in \{1, \dots, M\}\}$, frame size $T[k] = b_0 + \sum_{i=1}^M b_i x_i[k]$, and where we seek to minimize $(a_0 + \sum_{i=1}^M a_i \bar{x}_i)/(b_0 + \sum_{i=1}^M b_i \bar{x}_i)$ subject to $\sum_{i=1}^M c_{il} \bar{x}_i \leq d_l$ for all $l \in \{1, \dots, L\}$.

a) State the drift-plus-penalty ratio algorithm (37) for this, and conclude that:

$$\lim_{K \rightarrow \infty} \frac{a_0 + \sum_{i=1}^M a_i \bar{x}_i[K]}{b_0 + \sum_{i=1}^M b_i \bar{x}_i[K]} \leq ratio^{opt} + B/V, \quad \lim_{K \rightarrow \infty} \sum_{i=1}^M c_{il} \bar{x}_i[K] \leq d_l \quad \forall l \in \{1, \dots, L\}$$

for some finite constant B , where $ratio^{opt}$ is the optimal value of the objective function, and $\bar{x}_i[K] \triangleq \frac{1}{K} \sum_{k=0}^{K-1} x_i[k]$. Thus, the limiting time average satisfies the constraints and is within B/V from optimality. Your answer should solve for values $\phi_i[k]$ such that on frame k we choose $(x_1[k], \dots, x_M[k])$ over \mathcal{A} to minimize $\frac{\phi_0[k] + \sum_{i=1}^M \phi_i[k] x_i[k]}{b_0 + \sum_{i=1}^M b_i x_i[k]}$.

Note: In Appendix B it is shown that this minimization is performed as follows: Define $\mathcal{I} \triangleq \{i \in \{1, \dots, M\} | b_i = 0\}$ and $\mathcal{J} \triangleq \{j \in \{1, \dots, M\} | b_j > 0\}$. For all $i \in \mathcal{I}$, choose $x_i[k] = 0$ if $\phi_i[k] \geq 0$, and $x_i[k] = 1$ if $\phi_i[k] < 0$. Next, temporarily select $x_j[k] = 0$ for all $j \in \mathcal{J}$. Then rank order the indices $j \in \mathcal{J}$ from smallest to largest value of $\phi_j[k]/b_j$, and, using this order, greedily change $x_j[k]$ from 0 to 1 if it improves the solution. We stop when we reach the first index in the rank ordering that does not improve the solution.

b) Note that the case $b_0 = 1$ and $b_i = 0$ for $i \neq 0$ is a linear program. Give an explicit decision rule for each $x_i[k]$ in this case (the solution should be separable for each $i \in \{1, \dots, M\}$).

IV. REACTING TO RANDOMLY OBSERVED EVENTS

Consider a problem with general attributes $y_0[k], y_1[k], \dots, y_L[k]$ and frame size $T[k]$ for each frame k , as in Section III. However, now assume the controller observes a random event $\omega[k]$ at the beginning of each frame k , and this can influence attributes and frame sizes. The value of $\omega[k]$ can represent a vector of channel states and/or prices observed for frame k . Assume $\{\omega[k]\}_{k=0}^\infty$ is independent and identically distributed (i.i.d.) over frames. The controller chooses an action $\alpha[k] \in \mathcal{A}(\omega[k])$, where the action space $\mathcal{A}(\omega[k])$ possibly depends on the observed $\omega[k]$. Values of $(T[k], y_0[k], \dots, y_L[k])$ are conditionally independent of the past given the current $\omega[k]$ and $\alpha[k]$, with mean values:

$$\hat{y}_l(\omega[k], \alpha[k]) = \mathbb{E}[y_l[k] | \omega[k], \alpha[k]] \quad , \quad \hat{T}(\omega[k], \alpha[k]) = \mathbb{E}[T[k] | \omega[k], \alpha[k]]$$

The goal is to solve the following optimization problem:

$$\text{Minimize: } \bar{y}_0/\bar{T} \quad (48)$$

$$\text{Subject to: } \bar{y}_l/\bar{T} \leq c_l \quad \forall l \in \{1, \dots, L\} \quad (49)$$

$$\alpha[k] \in \mathcal{A}(\omega[k]) \quad \forall k \in \{0, 1, 2, \dots\} \quad (50)$$

As before, the constraints $\bar{y}_l \leq c_l \bar{T}$ are satisfied via virtual queues $Q_l[k]$ for $l \in \{1, \dots, L\}$:

$$Q_l[k+1] = \max[Q_l[k] + y_l[k] - c_l T[k], 0] \quad (51)$$

The random $\omega[k]$ observations make this problem more complex than those considered in previous sections of this paper. We present two different algorithms from [1][2].

Algorithm 1: Every frame k , observe $\mathbf{Q}[k]$ and $\omega[k]$ and choose $\alpha[k] \in \mathcal{A}(\omega[k])$ to minimize the following ratio of expectations:

$$\frac{\mathbb{E}\left[V\hat{y}_0(\omega[k], \alpha[k]) + \sum_{l=1}^L Q_l[k]\hat{y}_l(\omega[k], \alpha[k])|\mathbf{Q}[k]\right]}{\mathbb{E}\left[\hat{T}(\omega[k], \alpha[k])|\mathbf{Q}[k]\right]} \quad (52)$$

Then update the virtual queues via (51).

Algorithm 2: Define $\theta[0] = 0$, and define $\theta[k]$ for $k \in \{1, 2, 3, \dots\}$ as a running ratio of averages over past frames:

$$\theta[k] \triangleq \sum_{i=0}^{k-1} y_0[i] / \sum_{i=0}^{k-1} T[i] \quad (53)$$

Then every frame k , observe $\mathbf{Q}[k]$, $\theta[k]$, and $\omega[k]$, and choose $\alpha[k] \in \mathcal{A}(\omega[k])$ to minimize the following function:

$$V[\hat{y}_0(\omega[k], \alpha[k]) - \theta[k]\hat{T}(\omega[k], \alpha[k])] + \sum_{l=1}^L Q_l[k][\hat{y}_l(\omega[k], \alpha[k]) - c_l \hat{T}(\omega[k], \alpha[k])] \quad (54)$$

Then update the virtual queues via (51).

A. Comparison of Algorithms 1 and 2

Both algorithms are introduced and analyzed in [1][2], where they are shown to satisfy the desired constraints and yield an optimality gap of $O(1/V)$. Algorithm 1 can be analyzed in a manner similar to the proof of Theorem 1, and has the same tradeoff with V as given in that theorem. However, the ratio of expectations (52) is *not necessarily* minimized by observing $\omega[k]$ and choosing $\alpha[k] \in \mathcal{A}(\omega[k])$ to minimize the deterministic ratio given $\omega[k]$. In fact, the minimizing policy depends on the (typically unknown) probability distribution for $\omega[k]$. A more complex *bisection algorithm* is needed for implementation, as specified in [1][2].

Algorithm 2 is much simpler and involves a greedy selection of $\alpha[k] \in \mathcal{A}(\omega[k])$ based only on observation of $\omega[k]$, without requiring knowledge of the probability distribution for $\omega[k]$. However, its mathematical analysis does not yield as explicit information regarding convergence time as does Algorithm 1. Further, it requires a running average to be kept starting at frame 0, and hence may not be as adaptive when system statistics change. A more adaptive approximation of Algorithm 2 would define the average $\theta[k]$ over a moving window of some fixed number of frames, or would use an exponentially decaying average.

Both algorithms reduce to the following simplified *drift-plus-penalty* rule in the special case when the frame size $\hat{T}(\omega[k], \alpha[k])$ is a fixed constant T for all $\omega[k], \alpha[k]$: Every frame k , observe $\omega[k]$ and $Q[k]$ and choose $\alpha[k] \in \mathcal{A}(\omega[k])$ to minimize:

$$V\hat{y}_0(\omega[k], \alpha[k]) + \sum_{l=1}^L Q_l[k]\hat{y}_l(\omega[k], \alpha[k]) \quad (55)$$

Then update the virtual queues via (51). This special case algorithm was developed in [14] to treat systems with fixed size time slots.

A simulation comparison of the algorithms is given in [2]. The next subsection describes an application to energy-aware computation and transmission in a wireless smart phone. Exercise 7 considers opportunistic scheduling where wireless transmissions can be deferred by waiting for more desirable channels. Exercise 8 considers an example of price-aware energy consumption for a network server that can process computational tasks or outsource them to another server.

B. Efficient Computation and Transmission for a Wireless Smart Device

Consider a wireless smart device (such as a smart phone or sensor) that always has tasks to process. Each task involves a computation operation, followed by a transmission operation over a wireless channel. On each frame k , the device takes a new task and looks at its *meta-data* $\beta[k]$, being information that characterizes the task in terms of its computational and transmission requirements. Let d represent the time required to observe this meta-data. The device then chooses a *computational processing mode* $m[k] \in \mathcal{M}(\beta[k])$, where $\mathcal{M}(\beta[k])$ is the set of all mode options under $\beta[k]$. The mode $m[k]$ and meta-data $\beta[k]$ affect a computation time $D_{comp}[k]$, computation energy $e_{comp}[k]$, computation quality $q[k]$, and generate $A[k]$ bits for transmission over the channel. The expectations of $D_{comp}[k]$, $e_{comp}[k]$, $q[k]$ are:

$$\begin{aligned} \hat{D}_{comp}(\beta[k], m[k]) &= \mathbb{E}[D_{comp}[k]|\beta[k], m[k]] \\ \hat{e}_{comp}(\beta[k], m[k]) &= \mathbb{E}[e_{comp}[k]|\beta[k], m[k]] \\ \hat{q}(\beta[k], m[k]) &= \mathbb{E}[q[k]|\beta[k], m[k]] \end{aligned}$$

For example, in a wireless sensor, the mode $m[k]$ may represent a particular sensing task, where different tasks can have different qualities and thus incur different energies, times, and $A[k]$ bits for transmission. The full conditional distribution of $A[k]$, given $\beta[k]$, $m[k]$, will play a role in the transmission stage (rather than just its conditional expectation).

The $A[k]$ units of data must be transmitted over a wireless channel. Let $S[k]$ be the state of the channel on frame k , and assume $S[k]$ is constant for the duration of the frame. We choose a *transmission mode* $g[k] \in \mathcal{G}$, yielding a transmission time $D_{tran}[k]$ and transmission energy $e_{tran}[k]$ with expectations that depend on $S[k]$, $g[k]$, and $A[k]$. Define random event $\omega[k] = (\beta[k], S[k])$ and action $\alpha[k] = (m[k], g[k])$. We can then define expectation functions $\hat{D}_{tran}(\omega[k], \alpha[k])$ and $\hat{e}_{tran}(\omega[k], \alpha[k])$ by:

$$\hat{D}_{tran}(\omega[k], \alpha[k]) = \mathbb{E}[D_{tran}[k]|\omega[k], \alpha[k]] \quad , \quad \hat{e}_{tran}(\omega[k], \alpha[k]) = \mathbb{E}[e_{tran}[k]|\omega[k], \alpha[k]]$$

where the above expectations are defined via the *conditional distribution* associated with the number of bits $A[k]$ at the computation output, given the meta-data $\beta[k]$ and computation mode $m[k]$ selected in the computation phase (where $\beta[k]$ and $m[k]$ are included in the $\omega[k]$, $\alpha[k]$ information). The total frame size is thus $T[k] = d + D_{comp}[k] + D_{tran}[k]$.

The goal is to maximize frame processing quality per unit time \bar{q}/\bar{T} subject to a processing rate constraint of $1/\bar{T} \geq \lambda$, and subject to an average power constraint $(\bar{e}_{comp} + \bar{e}_{tran})/\bar{T} \leq P_{av}$ (where λ and P_{av} are given constants). This fits the general framework with observed random events $\omega[k] = (\beta[k], S[k])$, control actions $\alpha[k] = (m[k], g[k])$, and action space $\mathcal{A}(\omega[k]) = \mathcal{M}(\beta[k]) \times \mathcal{G}$. We can define $y_0[k] = -q[k]$, $y_1[k] = T[k] - 1/\lambda$, and $y_2[k] = e_{comp}[k] + e_{tran}[k] - P_{av}T[k]$, and solve the problem of minimizing \bar{y}_0/\bar{T} subject to $\bar{y}_1 \leq 0$ and $\bar{y}_2 \leq 0$. To do so, let $Q[k]$ and $Z[k]$ be virtual queues for the two constraints:

$$Q[k+1] = \max[Q[k] + T[k] - 1/\lambda, 0] \quad (56)$$

$$Z[k+1] = \max[Z[k] + e_{comp}[k] + e_{tran}[k] - P_{av}T[k], 0] \quad (57)$$

Using Algorithm 2, we define $\theta[0] = 0$ and $\theta[k]$ for $k \in \{1, 2, \dots\}$ by (53). The Algorithm 2 minimization (54) amounts to observing $(\beta[k], S[k])$, $Q[k]$, $Z[k]$, and $\theta[k]$ on each frame k and choosing $m[k] \in \mathcal{M}(\beta[k])$ and $g[k] \in \mathcal{G}$ to minimize:

$$\begin{aligned} & V[-\hat{q}(\beta[k], m[k]) - \theta[k](d + \hat{D}_{comp}(\beta[k], m[k]) + \hat{D}_{tran}(\omega[k], \alpha[k]))] \\ & + Q[k](d + \hat{D}_{comp}(\beta[k], m[k]) + \hat{D}_{tran}(\omega[k], \alpha[k]) - 1/\lambda) \\ & + Z[k](\hat{e}_{comp}(\beta[k], m[k]) + \hat{e}_{tran}(\omega[k], \alpha[k]) - P_{av}(d + \hat{D}_{comp}(\beta[k], m[k]) + \hat{D}_{tran}(\omega[k], \alpha[k]))) \end{aligned}$$

The computation and transmission operations are coupled and cannot be separated. This yields the following algorithm: Every frame k :

- Observe $\omega[k] = (\beta[k], S[k])$ and values $Q[k]$, $Z[k]$, $\theta[k]$. Then *jointly* choose action $m[k] \in \mathcal{M}(\beta[k])$ and $g[k] \in \mathcal{G}$, for a combined action $\alpha[k] = (m[k], g[k])$, to minimize:

$$\begin{aligned} & -V\hat{q}(\beta[k], m[k]) + \hat{D}_{comp}(\beta[k], m[k])[Q[k] - V\theta[k] - P_{av}Z[k] + Z[k]\hat{e}_{comp}(\beta[k], m[k]) \\ & + \hat{D}_{tran}(\omega[k], \alpha[k])Q[k] - V\theta[k] - P_{av}Z[k] + Z[k]\hat{e}_{tran}(\omega[k], \alpha[k])] \end{aligned}$$

- (Updates) Update $Q[k]$, $Z[k]$, $\theta[k]$ via (56), (57), and (53).

Exercise 9 shows the algorithm can be implemented without $\theta[k]$ if the goal is changed to maximize \bar{q} , rather than \bar{q}/\bar{T} . Further, the computation and transmission decisions can be *separated* if the system is modified so that the bits generated from computation are handed to a separate transmission layer for eventual transmission over the channel, rather than requiring transmission on the same frame, similar to the structure used in Exercise 8.

C. Exercises for Section IV

Exercise 7: (Energy-Efficient Opportunistic Scheduling [14]) Consider a wireless device that operates over fixed size time slots $k \in \{0, 1, 2, \dots\}$. Every slot k , new data of size $A[k]$ bits arrives and is added to a queue. The data must eventually be transmitted over a time-varying channel. At the beginning of every slot k , a controller observes the *channel state* $S[k]$ and allocates power $p[k]$ for transmission, enabling transmission of $\mu[k]$ bits, where $\mu[k] = \hat{\mu}(S[k], p[k])$ for some given function $\hat{\mu}(S, p)$. Assume $p[k]$ is chosen so that $0 \leq p[k] \leq P_{max}$ for some constant $P_{max} > 0$. We want to minimize average power \bar{p} subject to supporting all data, so that $\bar{A} \leq \bar{\mu}$. Treat this as a problem with all frames equal to 1 slot, observed random events $\omega[k] = (A[k], S[k])$, and actions $\alpha[k] = p[k] \in [0, P_{max}]$. Design an appropriate queue update and power allocation algorithm, using the policy structure (55).

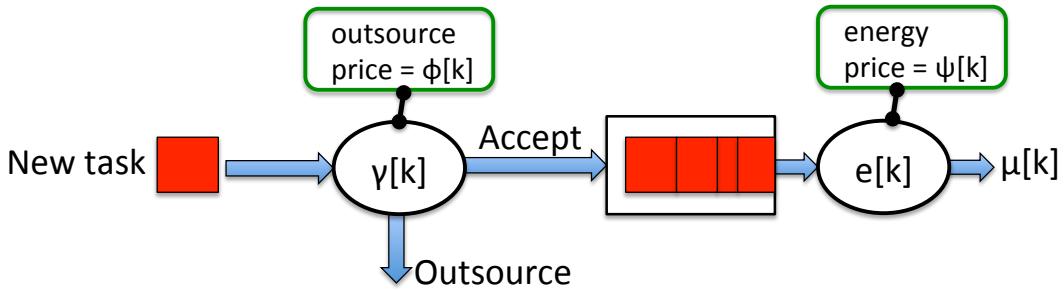


Fig. 10. The client/server system of Exercise 8, with decision variables $(\gamma[k], e[k])$ and observed prices $(\phi[k], \psi[k])$.

Exercise 8: (Energy Prices and Network Outsourcing) Consider a computer server that operates over fixed length time slots $k \in \{0, 1, 2, \dots\}$. Every slot k , a new task arrives and has size $S[k]$ (if no task arrives then $S[k] = 0$). The server decides to either accept the task, or outsource it to another server (see Fig. 10). Let $\gamma[k]$ be a binary decision variable that is 1 if the server accepts the task on slot k , and zero else. Define $A[k] = \gamma[k]S[k]$ as the total workload admitted on slot k , which is added to the queue of work to be done. Let $\phi[k]$ be the (possibly time-varying) cost per unit size for outsourcing, so the outsourcing cost is $c_{out}[k] = \phi[k](1 - \gamma[k])S[k]$. Every slot k , the server additionally decides to process some of its backlog by purchasing an amount of energy $e[k]$ at a per-unit energy price $\psi[k]$, serving $\mu[k] = \hat{\mu}(e[k])$ units of backlog with cost $c_{energy}[k] = \psi[k]e[k]$, where $\hat{\mu}(e)$

is some given function. Assume $e[k]$ is chosen in some interval $0 \leq e[k] \leq e_{max}$. The goal is to minimize time average cost $\bar{c}_{out} + \bar{c}_{energy}$ subject to supporting all tasks, so that $\bar{A} \leq \bar{\mu}$. Treat this as a problem with all frames equal to 1 slot, observed random events $\omega[k] = (S[k], \phi[k], \psi[k])$, actions $\alpha[k] = (\gamma[k], e[k])$, and action space $\mathcal{A}(\omega[k])$ being the set of all (γ, e) such that $\gamma \in \{0, 1\}$ and $0 \leq e \leq e_{max}$. Design an appropriate queue and state the dynamic algorithm (55) for this problem. Show that it can be implemented without knowledge of the probability distribution for $(S[k], \phi[k], \psi[k])$, and that the $\gamma[k]$ and $e[k]$ decisions are separable.

Exercise 9: Consider the same problem of Section IV-B, with the modified objective of maximizing \bar{q} subject to $1/\bar{T} \geq \lambda$ and $(\bar{e}_{comp} + \bar{e}_{tran})/\bar{T} \leq P_{av}$. Use the observation from Exercise 4 that this can be viewed as a problem with an “effective” fixed frame size equal to 1, and give an algorithm from the policy structure (55).

V. CONCLUSIONS

This paper presents a methodology for optimizing time averages in systems with variable length frames. Applications include energy and quality aware task scheduling in smart phones, cost effective energy management at computer servers, and more. The resulting algorithms are dynamic and often do not require knowledge of the probabilities that affect system events. While the performance theorem of this paper was stated under simple i.i.d. assumptions, the same algorithms are often provably robust to non-i.i.d. situations, including situations where the events are non-ergodic [1]. Simulations in Section III-D show examples of how such algorithms adapt to changes in the event probabilities. Exercises in this paper were included to help readers learn to design dynamic algorithms for their own optimization problems.

The solution technique of this paper uses the theory of optimization for renewal systems from [1][2], which applies to general problems. Performance for individual problems can often be improved using enhanced techniques, such as using place-holder backlog, exponential Lyapunov functions, LIFO scheduling, and ϵ -persistent service queues, as discussed in [1] and references therein. However, the drift-plus-penalty methodology described in this paper provides much of the insight needed for these more advanced techniques. It is also simple to work with and typically yields desirable solution quality and convergence time.

APPENDIX A — BOUNDED MOMENT CONVERGENCE THEOREM

This appendix provides a *bounded moment convergence theorem* that is often more convenient than the standard Lebesgue dominated convergence theorem (see, for example, [43] for the standard Lebesgue dominated convergence theorem). We are unaware of a statement and proof in the literature, so we give one here for completeness. Let $X(t)$ be a random process defined either over non-negative real numbers $t \geq 0$, or discrete time $t \in \{0, 1, 2, \dots\}$. Recall that $X(t)$ converges *in probability* to a constant x if for all $\epsilon > 0$ we have:

$$\lim_{t \rightarrow \infty} \Pr[|X(t) - x| > \epsilon] = 0$$

Theorem 2: Suppose there is a real number x such that $X(t)$ converges to x in probability. Further suppose there are finite constants $C > 0$, $\delta > 0$ such that for all t we have $\mathbb{E}[|X(t)|^{1+\delta}] \leq C$. Then $\lim_{t \rightarrow \infty} \mathbb{E}[X(t)] = x$.

Proof: Without loss of generality, assume $x = 0$ (else, we can define $Y(t) = X(t) - x$). Fix $\epsilon > 0$. By definition of $X(t)$ converging to 0 in probability, we have:

$$\lim_{t \rightarrow \infty} \Pr[|X(t)| > \epsilon] = 0 \quad (58)$$

Further, for all t we have $\mathbb{E}[X(t)] \leq \mathbb{E}[|X(t)|]$, and so:

$$\mathbb{E}[X(t)] \leq \epsilon + \mathbb{E}[|X(t)| \mid |X(t)| > \epsilon] \Pr[|X(t)| > \epsilon] \quad (59)$$

We want to show the final term in the right-hand-side above converges to 0 when $t \rightarrow \infty$. To this end, note that for all t we have:

$$\begin{aligned} C &\geq \mathbb{E}[|X(t)|^{1+\delta}] \\ &\geq \mathbb{E}[|X(t)|^{1+\delta} \mid |X(t)| > \epsilon] \Pr[|X(t)| > \epsilon] \\ &\geq \mathbb{E}[|X(t)| \mid |X(t)| > \epsilon]^{1+\delta} \Pr[|X(t)| > \epsilon] \end{aligned} \quad (60)$$

where (60) follows by Jensen's inequality applied to the conditional expectation of the function $f(|X(t)|)$, where $f(x) = x^{1+\delta}$ is convex over $x \geq 0$. Multiplying inequality (60) by $Pr[|X(t)| > \epsilon]^\delta$ yields:

$$CPr[|X(t)| > \epsilon]^\delta \geq (\mathbb{E}[|X(t)| \mid |X(t)| > \epsilon] Pr[|X(t)| > \epsilon])^{1+\delta} \geq 0$$

Taking a limit of the above as $t \rightarrow \infty$ and using (58) yields:

$$0 \geq \lim_{t \rightarrow \infty} (\mathbb{E}[|X(t)| \mid |X(t)| > \epsilon] Pr[|X(t)| > \epsilon])^{1+\delta} \geq 0$$

It follows that:

$$\lim_{t \rightarrow \infty} \mathbb{E}[|X(t)| \mid |X(t)| > \epsilon] Pr[|X(t)| > \epsilon] = 0$$

Using this equality and taking a \limsup of (59) yields $\limsup_{t \rightarrow \infty} \mathbb{E}[X(t)] \leq \epsilon$. This holds for all $\epsilon > 0$, and so $\limsup_{t \rightarrow \infty} \mathbb{E}[X(t)] \leq 0$. Similarly, it can be shown that $\liminf_{t \rightarrow \infty} \mathbb{E}[X(t)] \geq 0$. Thus, $\lim_{t \rightarrow \infty} \mathbb{E}[X(t)] = 0$. \square

Recall that convergence with probability 1 is stronger than convergence in probability, and so the above result also holds if $\lim_{t \rightarrow \infty} X(t) = x$ with probability 1. Theorem 2 can be applied to the case when $\lim_{K \rightarrow \infty} \frac{1}{K} \sum_{k=0}^{K-1} e[k] = \bar{e}$ with probability 1, for some finite constant \bar{e} , and when there is a constant C such that $\mathbb{E}[e[k]^2] \leq C$ for all k . Indeed, one can define $X(K) = \frac{1}{K} \sum_{k=0}^{K-1} e[k]$ for $K \in \{1, 2, 3, \dots\}$ and use the Cauchy-Schwartz inequality to show $\mathbb{E}[X(K)^2] \leq C$ for all $K \in \{1, 2, 3, \dots\}$, and so $\lim_{K \rightarrow \infty} \frac{1}{K} \sum_{k=0}^{K-1} \mathbb{E}[e[k]] = \bar{e}$.

APPENDIX B — THE DECISION RULE FOR LINEAR FRACTIONAL PROGRAMMING

This appendix proves optimality of the decision rule for online linear fractional programming given in Exercise 6. Specifically, suppose we have real numbers $\phi_0, \phi_1, \dots, \phi_M$, a positive real number $b_0 > 0$, and non-negative real numbers b_1, \dots, b_M . The goal is to solve the following problem:

$$\begin{aligned} \text{Minimize: } & \frac{\phi_0 + \sum_{i=1}^M \phi_i x_i}{b_0 + \sum_{i=1}^M b_i x_i} \\ \text{Subject to: } & 0 \leq x_i \leq 1 \quad \forall i \in \{1, \dots, M\} \end{aligned}$$

To this end, define disjoint sets \mathcal{I} and \mathcal{J} that partition the indices $\{1, \dots, M\}$ as follows:

$$\mathcal{I} \triangleq \{i \in \{1, \dots, M\} \mid b_i = 0\}, \quad \mathcal{J} \triangleq \{j \in \{1, \dots, M\} \mid b_j > 0\}$$

Then the objective function can be written:

$$\frac{\phi_0 + \sum_{i=1}^M \phi_i x_i}{b_0 + \sum_{i=1}^M b_i x_i} = \frac{\phi_0 + \sum_{i \in \mathcal{I}} \phi_i x_i}{b_0 + \sum_{j \in \mathcal{J}} b_j x_j} + \frac{\sum_{j \in \mathcal{J}} \phi_j x_j}{b_0 + \sum_{j \in \mathcal{J}} b_j x_j}$$

The denominator is always positive, and so regardless of the x_j values for $j \in \mathcal{J}$, the first term on the right-hand-side above is minimized by choosing x_i for all $i \in \mathcal{I}$ such that $x_i = 0$ if $\phi_i \geq 0$, and $x_i = 1$ if $\phi_i < 0$. Let x_i^* for $i \in \mathcal{I}$ denote these decisions. It remains only to choose optimal x_j values for $j \in \mathcal{J}$. Define $\gamma_0 \triangleq \phi_0 + \sum_{i \in \mathcal{I}} \phi_i x_i^*$. We want to solve the following:

$$\text{Minimize: } \frac{\gamma_0 + \sum_{j \in \mathcal{J}} \phi_j x_j}{b_0 + \sum_{j \in \mathcal{J}} b_j x_j} \quad (61)$$

$$\text{Subject to: } 0 \leq x_j \leq 1 \quad \forall j \in \mathcal{J} \quad (62)$$

Define val^* as the minimum value for the objective function in the problem (61)-(62).

Lemma 4: The following values $(x_j^*)_{j \in \mathcal{J}}$ solve the problem (61)-(62): For each $j \in \mathcal{J}$, we have:

$$x_j^* = \begin{cases} 0 & \text{if } \phi_j/b_j \geq val^* \\ 1 & \text{if } \phi_j/b_j < val^* \end{cases} \quad (63)$$

Proof: Let $(x_j)_{j \in \mathcal{J}}$ be an optimal solution to (61)-(62), so that $0 \leq x_j \leq 1$ for all $j \in \mathcal{J}$ and:

$$val^* = \frac{\gamma_0 + \sum_{j \in \mathcal{J}} \phi_j x_j}{b_0 + \sum_{j \in \mathcal{J}} b_j x_j}$$

Define c and d by:

$$c \triangleq \gamma_0 + \sum_{j \in \mathcal{J}} \phi_j x_j, \quad d \triangleq b_0 + \sum_{j \in \mathcal{J}} b_j x_j$$

so that $val^* = c/d$. Now take any $k \in \mathcal{J}$. It suffices to show we can change x_k to x_k^* , as defined in (63), without changing the ratio from its value val^* . Indeed, if this is true, then we can sequentially change each component x_j to x_j^* without changing the ratio, and so $(x_j^*)_{j \in \mathcal{J}}$ is also an optimal solution.

To this end, define val as the corresponding ratio with x_k replaced with x_k^* . We show that $val = val^*$. Clearly $val \geq val^*$ by definition of val^* as the minimum ratio. To show the opposite inequality, define $\delta_k^* = x_k^* - x_k$, and note that:

$$\begin{aligned}
val &= \frac{c + \phi_k \delta_k^*}{d + b_k \delta_k^*} \\
&= \frac{c}{d} + \frac{c + \phi_k \delta_k^*}{d + b_k \delta_k^*} - \frac{c}{d} \\
&= \frac{c}{d} + \frac{cd + \phi_k \delta_k^* d - cd - cb_k \delta_k^*}{d(d + b_k \delta_k^*)} \\
&= \frac{c}{d} + \delta_k^* [\phi_k/b_k - c/d] \frac{b_k}{d + b_k \delta_k^*} \\
&= val^* + \delta_k^* [\phi_k/b_k - val^*] \frac{b_k}{d + b_k \delta_k^*}
\end{aligned} \tag{64}$$

Note that $b_k/(d + b_k \delta_k^*) \geq 0$, and hence the second term on the right-hand-side of (64) is non-positive if $\delta_k^* [\phi_k/b_k - val^*] \leq 0$. In the case $\phi_k/b_k \geq val^*$, then by (63) we have $x_k^* = 0$. Thus, $\delta_k^* \leq 0$ and $\delta_k^* [\phi_k/b_k - val^*] \leq 0$. Thus, the second term in the right-hand-side of (64) is non-positive and so $val \leq val^*$. In the opposite case $\phi_k/b_k < val^*$, then $x_k^* = 1$ and so $\delta_k^* \geq 0$, $\delta_k^* [\phi_k/b_k - val^*] \leq 0$. Thus, the second term on the right-hand-side of (64) is again non-positive, so that $val \leq val^*$. \square

Now rank order the indices $j \in \mathcal{J}$ from smallest to largest value of ϕ_j/b_j , breaking ties arbitrarily. The solution of Lemma 4 has the form where $x_j^* = 1$ for the first n indices in this rank order, and $x_j^* = 0$ for the remaining indices, for some value $n \in \{0, 1, \dots, |\mathcal{J}|\}$ (where $|\mathcal{J}|$ is the size of set \mathcal{J}). Define val_n as the value of the ratio when $x_j = 1$ for the first n indices in the rank order, and $x_j = 0$ for the remaining indices in \mathcal{J} . Start with $val_0 = \gamma_0/b_0$. According to the rank ordering of \mathcal{J} , successively change x_j from 0 to 1 if it strictly decreases the ratio, computing val_1 , val_2 , and so on, until we reach the first index that does *not* improve the ratio. By an argument similar to (64), it is easy to see that this can only occur at an index j such that ϕ_j/b_j is greater than or equal to the current ratio, which means it is also greater than or equal to val^* (since by definition, val^* is less than or equal to any achievable ratio). Thus, all indices $i \in \mathcal{J}$ that come after j in the rank ordering also have $\phi_i/b_i \geq val^*$, and this greedy approach has thus arrived at the solution (63).

REFERENCES

- [1] M. J. Neely. *Stochastic Network Optimization with Application to Communication and Queueing Systems*. Morgan & Claypool, 2010.
- [2] M. J. Neely. Dynamic optimization and learning for renewal systems. *Proc. Asilomar Conf. on Signals, Systems, and Computers*, Nov. 2010.
- [3] E. Grochowski, R. Ronen, J. Shen, and H. Wang. Best of both latency and throughput. *Proc. IEEE Conf. on Computer Design (ICCD)*, pp. 236-243, October 2004.
- [4] M. Annaram, E. Grochowski, and J. Shen. Mitigating amdahl's law through epi throttling. *Proc. 32nd International Symposium on Computer Architecture (ISCA)*, pp. 298-309, June 2005.
- [5] R. Gallager. *Discrete Stochastic Processes*. Kluwer Academic Publishers, Boston, 1996.
- [6] S. Ross. *Introduction to Probability Models*. Academic Press, 8th edition, Dec. 2002.
- [7] S. Boyd and L. Vandenberghe. *Convex Optimization*. Cambridge University Press, 2004.
- [8] C. Li and M. J. Neely. Network utility maximization over partially observable markovian channels. *Arxiv Technical Report: arXiv:1008.3421v1*, Aug. 2010.
- [9] C. Li and M. J. Neely. Network utility maximization over partially observable markovian channels. *Proc. Intl. Symposium on Modeling and Optimization in Mobile, Ad Hoc, and Wireless Networks (WiOpt)*, May 2011.
- [10] L. Tassiulas and A. Ephremides. Stability properties of constrained queueing systems and scheduling policies for maximum throughput in multihop radio networks. *IEEE Transactions on Automatic Control*, vol. 37, no. 12, pp. 1936-1948, Dec. 1992.
- [11] L. Tassiulas and A. Ephremides. Dynamic server allocation to parallel queues with randomly varying connectivity. *IEEE Transactions on Information Theory*, vol. 39, no. 2, pp. 466-478, March 1993.
- [12] M. J. Neely. *Dynamic Power Allocation and Routing for Satellite and Wireless Networks with Time Varying Channels*. PhD thesis, Massachusetts Institute of Technology, LIDS, 2003.
- [13] L. Georgiadis, M. J. Neely, and L. Tassiulas. Resource allocation and cross-layer control in wireless networks. *Foundations and Trends in Networking*, vol. 1, no. 1, pp. 1-149, 2006.

- [14] M. J. Neely. Energy optimal control for time varying wireless networks. *IEEE Transactions on Information Theory*, vol. 52, no. 7, pp. 2915-2934, July 2006.
- [15] M. J. Neely. Queue stability and probability 1 convergence via lyapunov optimization. *Arxiv Technical Report, arXiv:1008.3519v2*, Oct. 2010.
- [16] F. Kelly. Charging and rate control for elastic traffic. *European Transactions on Telecommunications*, vol. 8, no. 1 pp. 33-37, Jan.-Feb. 1997.
- [17] J. Mo and J. Walrand. Fair end-to-end window-based congestion control. *IEEE/ACM Transactions on Networking*, vol. 8, no. 5, Oct. 2000.
- [18] A. Tang, J. Wang, and S. Low. Is fair allocation always inefficient. *Proc. IEEE INFOCOM*, March 2004.
- [19] W.-H. Wang, M. Palaniswami, and S. H. Low. Application-oriented flow control: Fundamentals, algorithms, and fairness. *IEEE/ACM Transactions on Networking*, vol. 14, no. 6, Dec. 2006.
- [20] I. Hou and P. R. Kumar. Utility maximization for delay constrained qos in wireless. *Proc. IEEE INFOCOM*, March 2010.
- [21] I. Hou, V. Borkar, and P. R. Kumar. A theory of qos for wireless. *Proc. IEEE INFOCOM*, April 2009.
- [22] F.P. Kelly, A.Maulloo, and D. Tan. Rate control for communication networks: Shadow prices, proportional fairness, and stability. *Journ. of the Operational Res. Society*, vol. 49, no. 3, pp. 237-252, March 1998.
- [23] S. H. Low. A duality model of TCP and queue management algorithms. *IEEE Trans. on Networking*, vol. 11, no. 4, pp. 525-536, August 2003.
- [24] L. Xiao, M. Johansson, and S. P. Boyd. Simultaneous routing and resource allocation via dual decomposition. *IEEE Transactions on Communications*, vol. 52, no. 7, pp. 1136-1144, July 2004.
- [25] X. Lin and N. B. Shroff. Joint rate control and scheduling in multihop wireless networks. *Proc. of 43rd IEEE Conf. on Decision and Control, Paradise Island, Bahamas*, Dec. 2004.
- [26] H. Kushner and P. Whiting. Asymptotic properties of proportional-fair sharing algorithms. *Proc. of 40th Annual Allerton Conf. on Communication, Control, and Computing*, 2002.
- [27] R. Agrawal and V. Subramanian. Optimality of certain channel aware scheduling policies. *Proc. 40th Annual Allerton Conference on Communication, Control, and Computing, Monticello, IL*, Oct. 2002.
- [28] A. Stolyar. Maximizing queueing network utility subject to stability: Greedy primal-dual algorithm. *Queueing Systems*, vol. 50, no. 4, pp. 401-457, 2005.
- [29] A. Eryilmaz and R. Srikant. Joint congestion control, routing, and mac for stability and fairness in wireless networks. *IEEE Journal on Selected Areas in Communications, Special Issue on Nonlinear Optimization of Communication Systems*, vol. 14, pp. 1514-1524, Aug. 2006.
- [30] Q. Li and R. Negi. Scheduling in wireless networks under uncertainties: A greedy primal-dual approach. *Arxiv Technical Report: arXiv:1001:2050v2*, June 2010.
- [31] M. J. Neely, E. Modiano, and C. Li. Fairness and optimal stochastic control for heterogeneous networks. *Proc. IEEE INFOCOM*, March 2005.
- [32] A. Eryilmaz and R. Srikant. Fair resource allocation in wireless networks using queue-length-based scheduling and congestion control. *Proc. IEEE INFOCOM*, March 2005.
- [33] A. Ribeiro and G. B. Giannakis. Separation principles in wireless networking. *IEEE Transactions on Information Theory*, vol. 56, no. 9, pp. 4488-4505, Sept. 2010.
- [34] M. J. Neely, A. S. Tehrani, and A. G. Dimakis. Efficient algorithms for renewable energy allocation to delay tolerant consumers. *1st IEEE International Conference on Smart Grid Communications*, Oct. 2010.
- [35] M. J. Neely. Opportunistic scheduling with worst case delay guarantees in single and multi-hop networks. *Proc. IEEE INFOCOM*, 2011.
- [36] M. J. Neely. Stock market trading via stochastic network optimization. *Proc. IEEE Conference on Decision and Control (CDC)*, Atlanta, GA, Dec. 2010.
- [37] M. J. Neely and L. Huang. Dynamic product assembly and inventory control for maximum profit. *Proc. IEEE Conf. on Decision and Control (CDC)*, Atlanta, GA, Dec. 2010.
- [38] L. Huang and M. J. Neely. Utility optimal scheduling in energy harvesting networks. *Proc. MobiHoc*, May 2011.
- [39] R. Urgaonkar, B. Urgaonkar, M. J. Neely, and A. Sivasubramaniam. Optimal power cost management using stored energy in data centers. *Proc. SIGMETRICS*, June 2011.
- [40] M. Gatzianas, L. Georgiadis, and L. Tassiulas. Control of wireless networks with rechargeable batteries. *IEEE Transactions on Wireless Communications*, vol. 9, no. 2, pp. 581-593, Feb. 2010.
- [41] L. Jiang and J. Walrand. *Scheduling and Congestion Control for Wireless and Processing Networks*. Morgan & Claypool, 2010.
- [42] L. Huang and M. J. Neely. Utility optimal scheduling in processing networks. *Proc. IFIP, Performance*, 2011.
- [43] D. Williams. *Probability with Martingales*. Cambridge Mathematical Textbooks, Cambridge University Press, 1991.